COMPUTER LABORATORY

HAEC OTIA STVDIA FOVENT

FIAT LVX

The University of Liverpool, Brownlow Hill. P.O. Box 147. Liverpool L69 3BX

PROCEEDINGS OF

NETWORKSHOP 2

UNIVERSITY OF LIVERPOOL

11-12 APRIL 1978

Network Unit of the Computer Board and Research Councils
University of Liverpool Computer Laboratory

CONTENTS

PREFACE

The programme for the second Networkshop was designed to ensure

(a) that participants were made aware of new developments in the definition of X25, high level protocols and communications hardware since the Glasgow Networkshop in September 1977;

and

(b) that the subject of Campus (or Local Area!) Networks was discussed in depth.

Partly because of lack of major progress in the protocols area, and partly because of the vast manpower effort devoted to the demonstration of EPSS in June 1978, there is a somewhat thin account of this subject in these Proceedings.

Considerable interest was stimulated by the sessions on Campus Networks and it is intended that this subject will be further discussed at a future Workshop at Cambridge in September 1978.

Thanks are due to the participants, speakers and particularly the Network Unit for the success of the Networkshop.

Dr. J.D. Rice
Liverpool University

REPORT FROM THE NETWORK UNIT

DR. R.A. ROSNER

REPORT FROM THE NETWORK UNIT

1. Network Hierarchy

Campus networks will be developed to give flexible interconnection among all sorts of devices, to allow for future expansion of computing facilities, to permit the distribution of resources as technology advances and to exploit the penetration of microprocessors.

Regional networks will be a pre-requisite for resource sharing among nearby universities or to concentrate traffic destined to escape from the region.

A national network will give access to the national exporting centres (funded by the Computer Board and the SRC), to special facilities such as the DAP and, eventually, to international centres.

Diagram 1 indicates the conceptual relationships between the various levels of the hierarchy. It does not imply any particular form of communications or switching. In practice, traffic may be carried by a combination of the Post Office's packet-switched network (PSS) and leased lines.

The main principles to be followed are:-

(a) rationalise the use of leased lines by concentrating traffic where feasible;

(b) by-pass levels of the hierarchy where traffic, economics or performance requirements justify;

(c) minimise the number of access methods and protocols to be implemented at any given installation;

(d) minimise the disruption during the transition from current arrangements to networking.

2. Post Office Developments

The Post Office has asked manufacturers to indicate what they could supply if they were to implement an X25 national network (PSS) by mid-1979. The sites to be covered are shown in diagram 2. Replies have been received and the Post Office is likely to invite the most feasible respondents to tender.

However, the matter is now very much in the arena of high politics and finance. A realistic date for the start of such a service is probably the first half of 1980.

Diagram 3 summarises relevant parts of the Post Office monopoly as it applies to our community. A detailed statement appears in Network News 3.

3. The Network Unit/NPL Packet-Switching Exchange

We have sent our specification to potential suppliers. Our main aim is to find equipment which is highly modular and could, within a single upwards-compatible technology, satisfy both campus and regional switching requirements. Three possible suppliers have given interesting responses - two are based on multi-microprocessor solutions while the third is constructed around a conventional minicomputer. Costs vary but a pure packet-switch driving about 6 HDLC lines at 48 kb would cost ∿ £40k. The PAD function is expensive and a campus switch might therefore be about £50k.

Several regions have expressed interest in the PSE and discussions are taking place with the potential suppliers. It is hoped that comments on the specification from the community can be incorporated into the final version once a supplier is chosen.

4. Regional Switching and the Post Office

The Post Office is interested in supplying regional switches. Its current estimate is that this would cost about £80k per annum to rent (independent of traffic volume). The figure, on the Post Office's own admission, is based on the only technology for which evidence was available to them. It is therefore not to be taken too literally since a considerably reduced quote could be expected from the use of more modern technology.

The Post Office monopoly is obviously of vital importance since, if the final cost of regional switch rental is too high, regional networking might prove economically unjustifiable. However, if the price is reasonable, Post Office provision would have operational advantages and would define the absolute standards of the interfaces as well as guaranteeing compatibility with PSS.

The unresolved issues are numerous. No realistic quotation is yet available. There are no timescales for the availability of kit nor even a chosen supplier.

Decisions are intimately linked with the supply of PSS. Ideally, the technology of PSS would be downwards extensible in price and performance to allow the same basic equipment to be used for regional switching. If this is not the case, the Post Office might agree to rent out kit conforming to our specifications and obtained from our chosen supplier. (This solution would have cost implications since the Post Office would have to maintain expertise in two different technologies).

Campus networks are dealt with elsewhere in these proceedings.

5. Study Group 4

The Post Office runs a group with membership from the potential user community to look at various aspects of a possible X25 PSS.

Study Group 4A is concerned with technical issues. It has explored HDLC in some detail and produced state diagrams. Various ambiguities and inconsistencies in X25 have been uncovered and areas where choice or options are possible have been identified. It was originally intended to produce a PSS technical guide but, given the uncertainty of the technology to be chosen, this is felt to be impossible. The main aim is to maintain ongoing links between technical experts inside and outside the Post Office.

Study Group 4B is seen as the basis of a PSS User Liaison Group and is discussing desirable facilities of a future network. Various application areas are under study including private networks connected to PSS, electronic mail, banking, databases and bureaux. Tariffs are also being discussed and the Post Office is aware that users' plans cannot proceed until charges are known. However, these are unlikely to be announced until 1979.

Diagram 4 indicates some possible tariffs for PSS and is based on a combination of figures published for EPSS and inspired guesswork. Diagram 5 is a very crude comparison of costs for leased lines against the use of PSS for current levels of traffic. It is assumed that leased line costs (which have remained unaltered for a long period of high inflation) will rise by about 30% in the near future. Of course, the advent of PSS will have a great influence on modes of working so that some points of diagram 5 currently above the break-even lines could be forced to fall below them.

6. High-Level Protocols

The main event since the last Networkshop was the publication of the File Transfer Protocol which is now being implemented in many places.

There has also been significant activity in the standards organisations both nationally (BSI) and internationally (ISO). Diagram 6 shows the corresponding specialist working groups and their chairmen.

It is widely realised that the various activities taking place all over the country on high-level protocols need to be coordinated if practical results are to be forthcoming on a reasonable timescale. The phases identified for each task are

- specification of protocol
- refinement
- implementation on several machines
- modification in the light of experience
- submission to standards bodies.

Such an exercise must involve university and polytechnic computer centres and computer science departments, Research Council laboratories, government research establishments, software houses and computer manufacturers. It is suggested that the Department of Industry sponsor the work (through its Computer Systems and Electronics Requirements Board).

7. Machine Ranges and Networking Standards

A letter has been sent to all computer manufacturers by the Computer Board. It indicates that all computers delivered after 1 January 1981 will be expected to conform to a set of standards including X25, the file transfer protocol and such additional protocols as emerge within the coming year. Clauses embodying these requirements will be included in the specification of future systems and tenders will be judged against the criteria laid down.

A small group from university 2900 sites is working with ICL on the specification of networking requirements for ICL-supported operating systems. It is possible that this could lead to a joint development exercise.

The SRC is committed to the integration of its two networks: the Interactive Computing Facility based on DEC 10's at Edinburgh and Manchester and the main

SRC network for the IBM machines at Rutherford and Daresbury. The main network is being converted from an EPSS-like protocol to X25. A contract has been placed with Hatfield Polytechnic to study methods for connecting DEC 10's to an X25 network together with implementation of non-proprietary higher-level protocols.

The Network Unit is keen to encourage the formation of network groups among users of other machine ranges. Ideally, the aim should be to coordinate (in collaboration with the manufacturer) the development of the required software for future standard networking on each type of computer and for each manufacturer-supported operating system.

8. The Future

With the impending termination of the Network Unit's life in October, attention is being devoted to the tasks which lie beyond that date. It is assumed that the Computer Board and the Research Councils will agree, at the highest levels, that the sharing of various facilities particularly for communications represents a rational use of resources. Another prerequisite for future planning is the formulation of a more systematic policy for the recurrent funding of computing in a networking environment. The Unit has already generated significant input to the Board's deliberations on this issue.

With these assumptions, there appear to be four main areas of activity:- the establishment of a joint network policy by or on behalf of the funding bodies, the preparation and approval of technical plans to carry the policy into effect, the implementation of the plans and, in parallel with all these, liaison and coordination among all those concerned with communications throughout the community.

Elaboration of the details of each activity and the search for mechanisms by which the aims could be achieved will constitute an important task in the remaining months.

Roland Rosner
Network Unit of the Computer Board
and Research Councils
21 April 1978

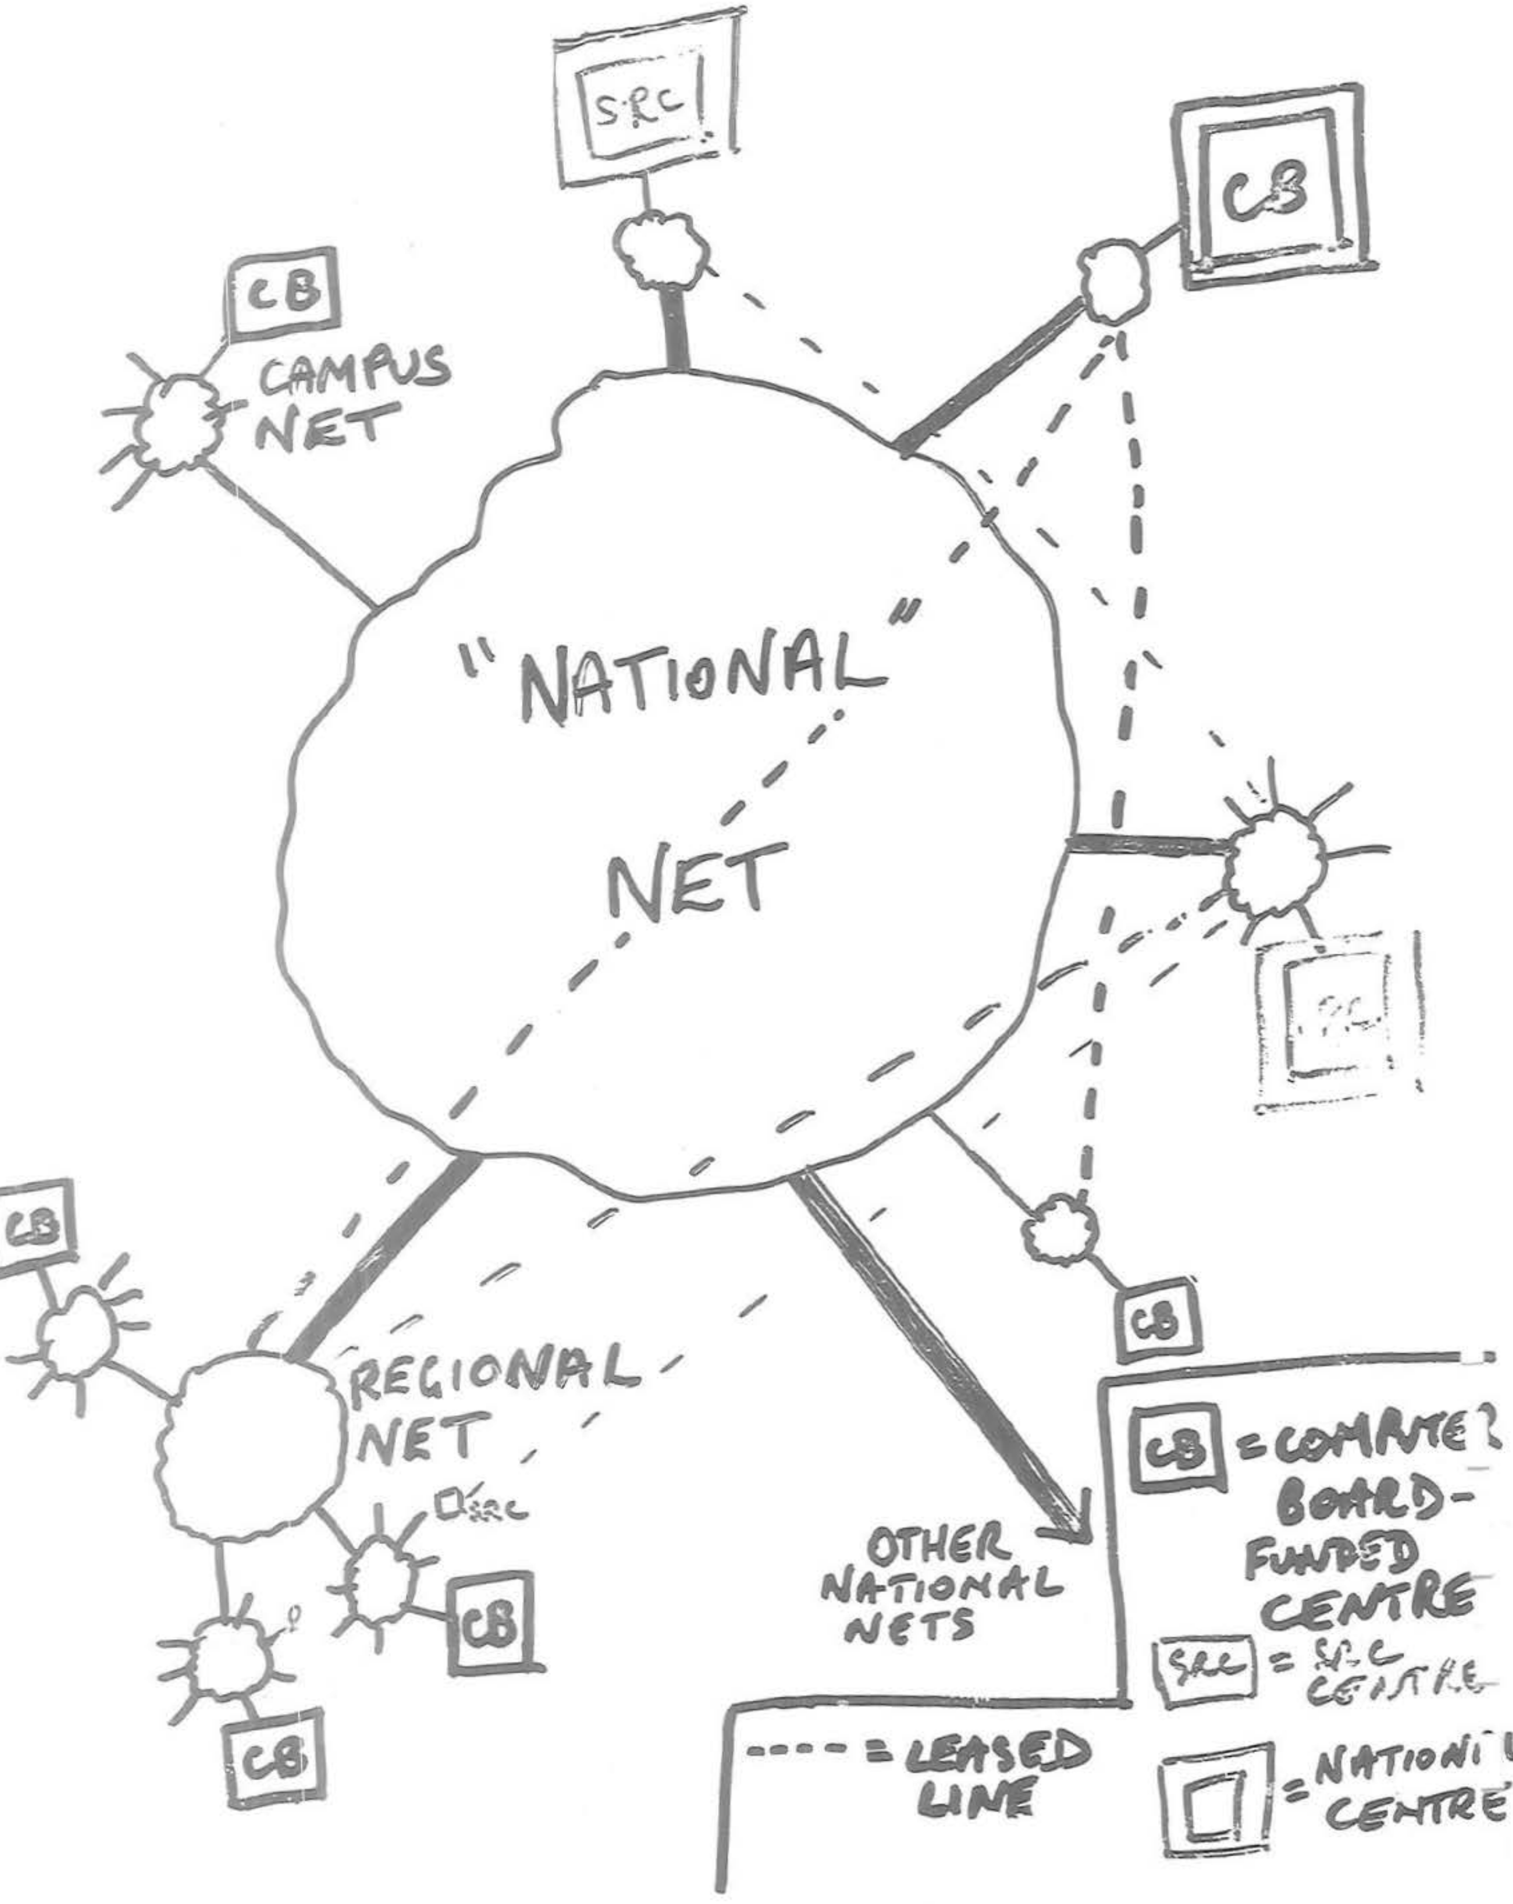

DIAGRAM 1 NETWORK HIERARCHY

POST OFFICE PACKET-DATA SERVICE

1980 COVERAGE

MAIN SWITCHING CENTRES	PACKET*		CHARACTER†
	UP TO 9.6	48K BIT/S	1200 BIT/S
LONDON	✓	✓	✓
BIRMINGHAM	✓	✓	✓
MANCHESTER	✓	✓	✓
ACCESS POINTS			
LEEDS	✓		✓
BRISTOL	✓		✓
CAMBRIDGE	✓		✓
READING	✓		✓
GLASGOW	✓		✓
EDINBURGH	✓		✓

†LEASED CIRCUIT OR DIAL-UP

*LEASED CIRCUIT ONLY

DIAGRAM 2

PO MONOPOLY

- PO HAS EXCLUSIVE PRIVILEGE OF SWITCHING DATA BETWEEN PARTIES BELONGING TO DIFFERENT LEGAL ENTITIES
- PO CAN GRANT LICENCES TO DO SUCH SWITCHING WHENEVER IT FEELS LIKE IT
- SWITCHING ENTIRELY WITHIN A CAMPUS REQUIRES NO LICENCE
- PO DOES NOT INTEND TO LICENCE PRIVATE REGIONAL SWITCHING (EXCEPT FOR EXISTING NETWORKS)
- PO EXPECTS TO BE ABLE TO OFFER REGIONAL SWITCHING FACILITIES ON A RENTAL BASIS (TRAFFIC-INDEPENDENT)
- SRC NETWORK REQUIRES LICENCE (UNIVERSITY USERS ARE NOT "SERVANTS" OF SRC WITHIN THE MEANING OF THE PO ACT)
- NATIONAL CENTRES AND SRC MACHINES CAN BE CONNECTED DIRECTLY TO PO-SUPPLIED REGIONAL SWITCHES

DIAGRAM 3

POST OFFICE PACKET-DATA SERVICE

Possible Tariffs

Volume	-	35 p per 1000 packets*
Call charge	-	£1 per hour
Port rental	-	£300-1000 per annum†

Multiple-access charge	£10 per channel per annum
Network user identity (Dial-up)	£20 per identity per annum

*1 packet = 128 bytes

†includes in-station modem and line up to 20 km

DIAGRAM 4

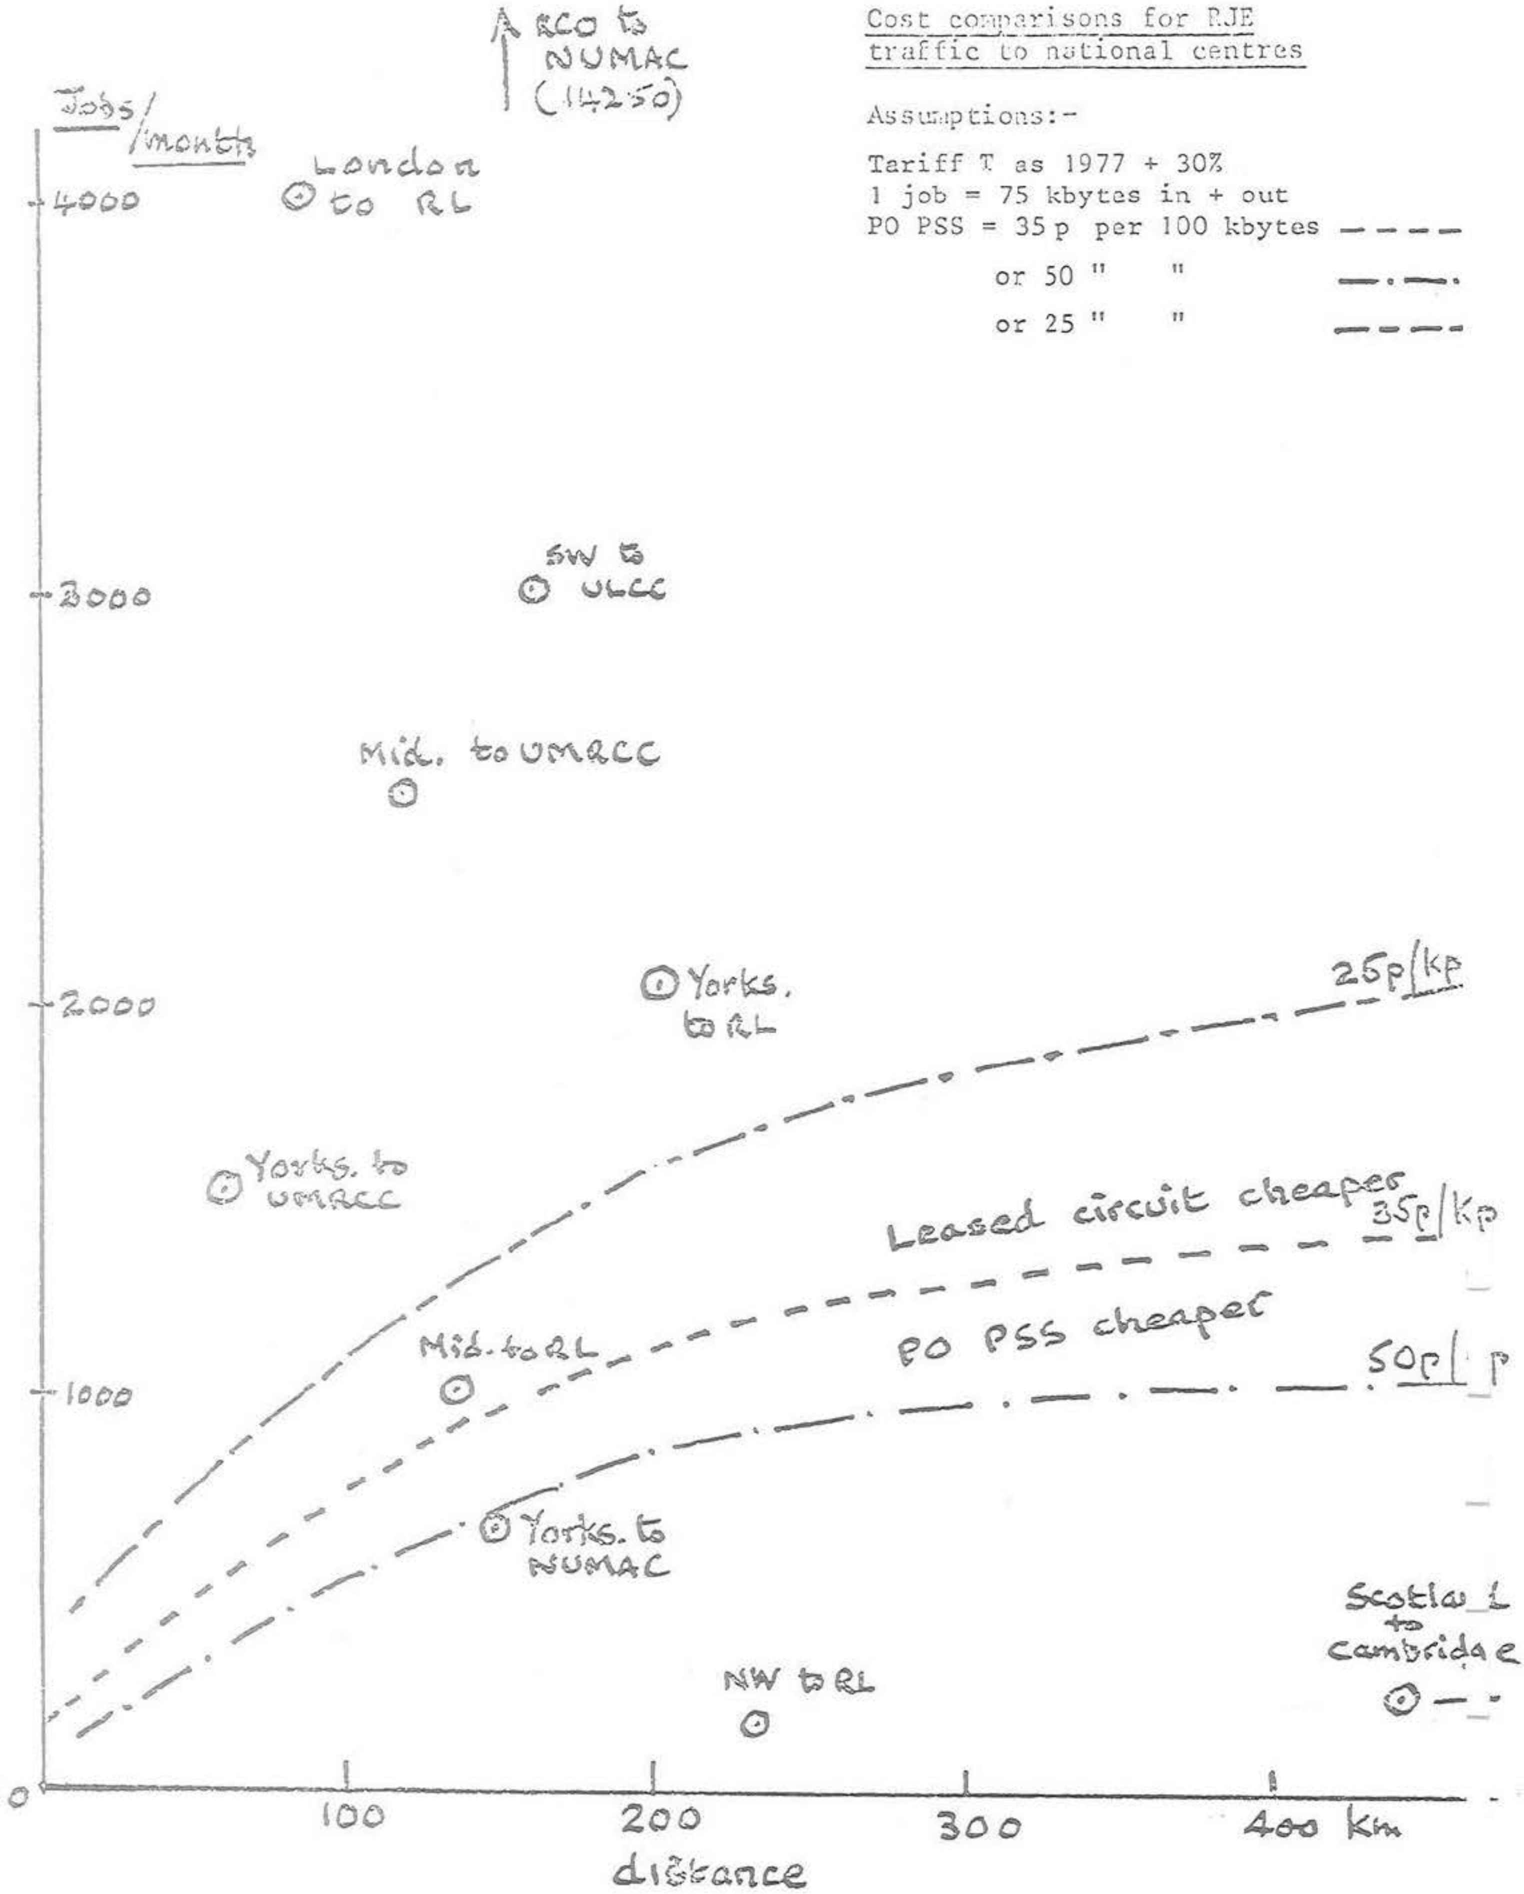
Cost comparisons for RJE traffic to national centres
Assumptions:-
Tariff T as 1977 + 30%
1 job = 75 kbytes in + out
PO PSS = 35 p per 100 kbytes
or 50 " "
or 25 " "
Jobs/month
4000
3000
2000
1000
0
100
200
300
400 Km
distance
RCO to NUMAC (14250)
London to RL
SW to ULCC
Mid. to UMRCC
Yorks. to RL
Yorks. to UMRCC
Mid. to RL
Yorks. to NUMAC
NW to RL
Scotland to Cambridge
25p/Kp
35p/Kp
50p/Kp
Leased circuit cheaper
PO PSS cheaper

DIAGRAM 5

OPEN SYSTEM INTERCONNECTION

BSI DPS/20 WG			ISO TC97/SC16 WG	
1	MODELS FOR NETWORK ARCHITECTURE STRUCTURE A. CHANDLER	→	1	ARCHITECTURE H. ZIMMERMANN (FRANCE)
2	TASK ACTIVATION AND INTER-TASK DIALOGUE A. LANGSFORD			
	TRANSPORT SERVICE FUNCTIONS K. HEARD	→	3	TRANSPORT SERVICE (USA)
4	VIRTUAL TERMINAL PROTOCOLS C. SOLOMONIDES	→	2	USERS OF TRANSPORT SERVICE A. LANGSFORD (U.K)
5	INFORMATION PRESENTATION R. ROSNER			
6	FILE TRANSFER AND MANAGEMENT P. LININGTON	→	2	

DIAGRAM 6

REPORT ON X25

DR. C. SOLOMONIDES, NPL

REPORT ON X25

Level 2

The LAP changes, to bring the procedures in line with ISO HDLC, and the introduction of LAP B have now been approved by CCITT in a postal vote. This gives the modified document the status of provisional recommendation.

Level 3

In October 1977 the final meeting of the editor's group, before the full study group meeting in April 78, took place in Geneva. A modified packet level goes to the study group for consideration. This includes changes to Sections three to five of the X25 document. Section 3 and 4 changes are mainly editorial to clarify the interface definition with some significant improvements and extensions. These include the addition of an eight bit diagnostic code in Restart and Clear packets as well as the Reset packet in which it was already present. The Canadian representatives at the meeting pressed for a unified definition of diagnostic codes for all three operations. Actions initiated by the DTE may make use of this diagnostic code field to carry user information as to the reason of not accepting or resetting a call. User initiated actions are distinguished from network ones by the cause field coding.

Packet formats for extended numbering have been added. This introduces new types for packets which contain level 3 sequence numbers.

Work has started on optional user facilities which are defined in Section 5 of the recommendation. The closed user group facility has now been introduced. There was no clear agreement on the exact nature and mechanisms of this facility with Japan puting forward proposals for a complex dynamic service. On this as well as other areas the Japanese are taking a somewhat iconoclastic view of the interface.

Further discussion on throughput classes showed some confussion in CCITT on whether packet and/or window size at level 3 have anything to do with throughput attainable by a virtual circuit. Throughput classes will guarantee a grade of service to the user. This discussion led to the concept of a default throughput class for all virtual circuits. Nevertheless it was not clear how this is achieved.

The subject of negotiation of optional user facilities was discussed and the possibility of extending the formats of the Call Actepted and Call Conected packets to include a facilities field for this purpose was considered.

Further Services.

Three further services are now under discussion in CCITT. These are:

1) Multiple line DTE/DCE interface. The French favour the X72 proposed solution of defining a multiline HDLC which preserves the sequential character of the protocol. Japan has other ideas.

2) A Datagram service which can be accessed via X25 but forms an independent network facility and can be used through interfaces other than X25.

3) Frame mode DTE to provide a simple single call interface when the multiplexing function of X25 is not required.

C M SOLOMONIDES

HIGH LEVEL PROTOCOLS

The Transport Service

Dr. K.S. Heard, AERE Harwell

The Network Independent File Transfer Protocol

Dr. P.F. Linington, Cambridge.

THE TRANSPORT SERVICE

1. Introduction

A transport service is intended to provide the means by which all higher levels of data processing systems may interact within heterogeneous computer communications networks.

The economic advantages of standardised access to a general inter-process communications service, through the use of a clearly defined functional interface, have become more apparent as the range and sophistication of communicating entities has increased. The provision of a standard transport service interface can isolate individual program applications from the details of any particular communications facility, and reduce the effort and complexity involved in implementation of all higher level protocols and processes.

The definition of a network independent set of functional transport service capabilities also provides an evolutionary solution to changes in network technology. Internal transport service procedures are made responsible for the detailed management of any specific lower level protocols that are required to improve on the facilities offered by any given underlying network and to provide the necessary degree of end-to-end control. The network independent interface provided by the transport service also provides a convenient way to interconnect different communication systems when this is considered desirable.

2. Transport Service Functions

A functional definition of the transport service can be derived from considerations of the requirements for general process interaction. A mapping of this functional capability onto a given communications mechanism is achieved through the specification and use of lower level interfaces and protocols that are matched to the detailed properties of the supporting systems.

The essence of any transport service can be stated as the provision of general purpose, secure, information transfer channels between pairs of processes. These channels should be accessible in a network independent manner for economic reasons.

The transport service should provide the following features:-

- process to process mapping
- assurance of grade of service
- transfer of data between processes
- process synchronisation
- end-to-end data flow control
- end-to-end error detection and control
- uniform access to services independent of transport medium

This functional definition of the transport service is derived in a 'top-down' approach in which the desired capabilities are seen through an application level interface to the transport service. How these functions are achieved is a matter for the transport service implementation.

A layered approach to both the definition and provision of the functional transport capability does however allow increasing levels of service to be added in a compatible manner. It also allows any individual layer below the service interface to be replaced without causing any functional change at the higher level. Figure 1 shows a layered structure for a transport service based on an X25 communications system. Note that each layer observes its own service protocol (shown by horizontal links). The protocols for levels 1-3 may cascade through several intermediate systems link by link. The upper layer (layer 4) of the transport service achieves end-to-end transport service functions, such as error and flow control, through maintenance of some protocol that concerns only the corresponding transport services.

The details of the end-to-end transport protocol and the supporting procedures will depend on the combination of the actual grade of service offered to the application processes and the facilities available from the lower service levels. A layered structure within the end-to-end transport layer is foreseen to provide alternative grades of service that are best suited to the needs of individual application processes, particularly with regard to response, throughput and cost.

TRANSPORT SERVICE - STRUCTURE

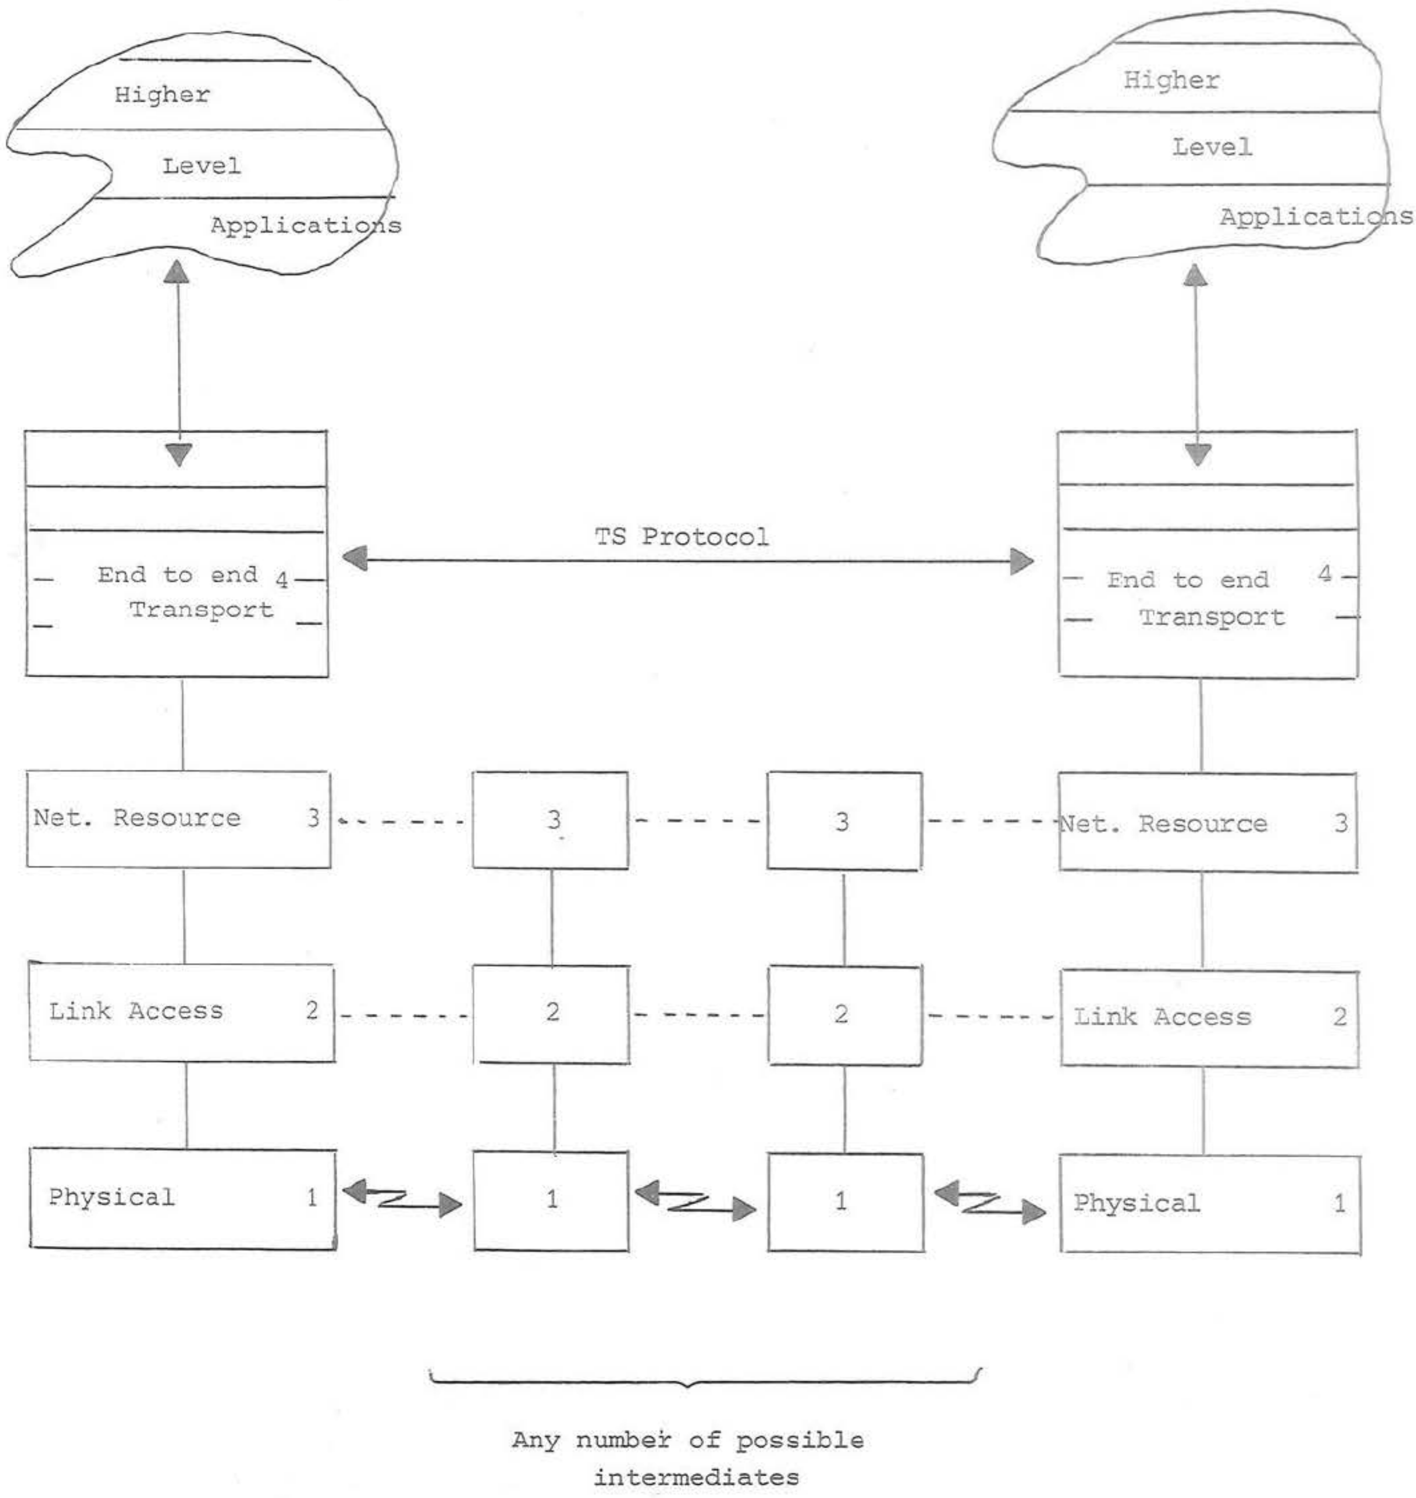

FIGURE 1

3. <u>Transport service primitives</u>

The transport service functions can be made accessible to higher level (application) processes through the use of primitive operators that penetrate the transport service interface. The necessary primitives can be sub-divided into four distinct categories corresponding to the application initiated functional requirements:-

- service quality control
- channel establishment and destruction
- transparent data transfer
- process synchronisation.

Selection and control of the grade of service will (in principle) involve negotiation between each application level process and its local transport service and may involve negotiation between corresponding transport service stations. The following incomplete list of parameters must be agreed with the transport service for proper management of service quality.

service quality control:

- routing information
- tolerable error rate
- message sequencing requirements
- flow control strategy
- throughput expectations
- tariff implications.

Some mechanisms must be provided to dynamically establish and subsequently break communications channels between processes, at least under the assumption of a need to conserve limited resources. The properties of the inter-process channel will generally be a subject for negotiation between the two (potentially) corresponding parties. Considerations of resource allocation in practical systems also suggest a real need for some method whereby an application process can indicate a willingness to consider unsolicited proposals to set up a communication channel.

Communication channel establishment and destruction:

. connect	–	propose channel attributes
. accept	–	confirm (modified) attributes
. reject	–	refuse to establish proposed channel
. disconnect	–	break an established channel
. listen	}	enable/disable local process sensitivity
. ignore		to incoming channel proposals

Application level information may be transferred over the selected communication channel through use of two simple primitives designed to pass data transparently across the transport service interface. Transfers are considered to concern finite strings of data. No other structure is attributed to the data by the transport service. The question of whether the end of a string indicates the end of a message (sometimes called a unit of interaction) has implications for flow control, response times, communications tariff charges and local transport service and/or process data buffering. It is a topic for further study.

data transfer:

. putdata	-	send data over a defined channel
. getdata	-	extend commitment to receive data
. killput	}	withdraw outstanding transfer requests
. killget		and modify flow control credits if relevant.

Synchronisation of two otherwise independent communicating processes may be achieved through the establishment and/or destruction of data transfer channels alone, or through a normal controlled transfer of data over these channels. Two other mechanisms are foreseen to cater for abnormal (error) situations in which a process may seek to alter the current chain of actions, or to recover from a detected loss of data. Process attention signalling should occur 'out of band' with respect to the normal data streams (which may be blocked) and may, or may not, destroy any data actually in transit according to the mechanism selected. Process sensitivity to signalling should be manipulated in the same way as for incoming channel proposals.

process synchronisation

. interrupt	-	non destructive process attention signalling
. reset	-	overall system re-initialisation
. enable sync	}	enable/disable process recognition of
. disable sync		received synchronisation signals.

4. <u>Current Working Group Activity</u>

There are a number of working groups concerned with definitions of a general transport service in both a national and international context:

- UKPO EPSS/SG3
- BSI DPS20/WG3
- IFIP INWG/TS
- ISO TC97/SC16/WG3

EPSS/SG3 has been particularly concerned withspecification of a network independent interface to a transport service to satisfy real needs for economic implementation on a publicly accessible network. Experience gained through implementation of the original Bridging Protocol has already been fed back into a greater understanding of practical requirements. This experience should be of benefit to both the user community and the Post Office, particularly since SG3 has now extended its scope of activities to consider the details of a transport protocol required above an X25 based network.

BSI DPS20/WG3 is a working group set up under the BSI committee concerned with all aspects of the problems to be encountered in open system interconnection. WG3 is attempting to define the functions and grades of servic required of a general inter-process communication facility, and to consider the transport level protocols which will be required to provide and control those functions on an end-to-end basis. DPS20/WG3 is specifically charged to map any proposals for a network independent transport service onto an X25 based network. Any work devoted to the preparation of prospective service standard must necessarily take into account the efforts of other concerned national and international bodies. DPS20 has established useful contact wiht users of various existing networks, with manufacturers, PTTs, IFIP and ISO.

IFIP INWG/TS have traditionally laid emphasis on the definition of an end-to-end protocol to be used with an insecure communications support service based on the datagram concept. Provision of a functional interface specificati was determined from consideration of the properties of the transport protocol rather than higher level considerations of application process requirements. Direct use of the transport protocol in a virtual call based environment may incur severe traffic penalties (in terms of response time and tariff charges). Recent attempts to modify the details of the protocol for use with more secure subnetworks are not entirely satisfactory.

ISO TC97/SC16/WG3 is essentially the international parallel to DPS20/WG3. The ISO transport service group is asked to identify the functions required by higher levels and to consider the structure and standardisation of all service levels below this functional interface.

5. A transport service above X25

An X25 based communications network will offer facilities corresponding to level 3 in Figure 1. The virtual call facility will offer:

- address selectable inter-process communication channels
- transfer of unlimited length data streams
- packet error detection and control (to a level of 1 in 10^8 ?)
- sequential packet delivery
- data flow control on a link by link basis
- out of band signalling (for process synchronisation).

Provision of general application level access to an X25 based transport service presents a number of immediate problems. Most of these problems are concerned with end-to-end control (level 4), which is not currently provided in X25. Other problems concern the agreement amongst all potentially interacting parties (users and PTTs) on the specification of addressing and service facility fields to allow commercial and technical decisions to be made concerning the interconnection and use of multi-level subscriber host computers and private networks.

The most immediate problems to be solved before a standardised X25 based transport service can be constructed seem to be:

- Subaddressing requirement for multi-level subscriber hosts and network gateways
- detailed specification of transport facilities required (throughput, charging, etc.)
- end-to-end data flow control and 'message' delimiters
- end-to-end error rate control (and improvement?)
- service grade monitoring and failure reporting
- controlled channel termination (without loss of data)

As topics for further study one could also list:

- multiplexed usage of an X25 virtual call channel
- broadcast utilisation of a group of virtual call channels
- tariff charge optimisation
- added value services such as encryption, dynamic service location

The need to integrate existing network address schemes into the provisions of an X25 based facility can be taken as an example of current work.

The X25 Recommendation foresees the use of 15 digit source and destination address fields to completely specify the routing information required to establish a communication channel between a pair of processes. This can be compared with the 12 digit international telephony scheme which enabl[es] one subscriber to contact any other anywhere in the world. But it does require a globally unified address scheme. Non-conformism is rewarded by exclusion from the accessible domain of the 'standard' network.

A universal format for the X25 address field has not yet been defined. It may be hierarchic, zonal or global. Conventional computer philosophy suggests that a hierarchic address scheme would be the most economic and convenient; such an address could be structured as:

network.switch.subscriber.local process.port

Unfortunately a 15 digit address field is unlikely to permit a general hierarchic address format. Any requirement to cascade addresses for the purpose of traversing network gateways will obviously require some more extended method for transfer of the necessary information.

Furthermore, it may be desirable to indicate specific routing information in addition to the source and final destination addresses. The ability to specify the route taken to establish a call channel with a given destination will often be necessary in the case of incompletely connected disjoint networks, and may be highly desirable from a security or service quality point of view.

Several alternative address schemes have been proposed, each with its own advantages and disadvantages:

user data field $\rightarrow$ 24 digits

The addresses could be extended into the user data field. Although this field is defined to be up to 16 octets long, only 12 octets are now available for use. This would offer an additional 24 address digits - enough to permit local process port identification for both source and destination, but not adequate for general gateway addresses.

facilities field → 62 digits

Use of the facilities field would certainly offer a significant expansion for the addressing capabilities. It is however very likely that this field will be policed by the communications subnetwork and therefore will not be available for anything except bona fide facility selection.

data phase packet → ∞ digits

Any subaddressing information can be carried in data packets transferred when the primary call has been established by normal X25 address field selection of a given target subscriber. This scheme imposes no limit whatsoever on the complexity of the total addressing information, but does demand a cascade commitment on each gateway involved in setting up the required communications channel. This may have serious consequences in terms of call set up response, resource allocation and tariff charges. The scheme would require adherence to a standard convention at the subscriber transport service level, but would not in any way impact on the operation of an X25 subnetwork. Furthermore, the use of subaddressing information continued within the data phase offers a scheme that can be applied to permanent virtual circuits (pvc) in exactly the same way as a switched virtual circuits (svc).

KENNETH S. HEARD

C.S.S.D.
H7.12
AERE Harwell

5th May 1978

The File Transfer Protocol

The ability to transfer programs and data between computers is essential for the use of distributed computing facilities. However, present day operating systems differ considerably in the techniques they use to handle and describe information. A standard method must be agreed for the description of files and of the operations to be performed on them, even if the use of a uniform Transport Service has removed communication specific problems.

A solution is proposed in the File Transfer Protocol produced by the High Level Protocol Group (formerly EPSS Study Group 2). This defines a Standard Conceptual Filestore, in terms of which the desired actions can be described. To use the protocol with an existing system, a mapping is made between the standard model and the particular environment; since the detail of the mapping does not effect the protocol exchanges, it is entirely local to the particular system. The same standard model can be used for communication between filing systems of differing sophistication, real devices or job processing facilities, since an appropriate mapping can be set out for each of these.

The protocol includes powerful facilities for identification of the required object, description of its properties and the control of the movement of its contents. This does not imply, however, that the use of the protocol implies a high implementational or operational cost. Indeed, simple needs can be met with very few control exchanges. The initial exchange that identifies the object under discussion is structured so that only those properties which are essential in the particular instance need be recognised; superfluous or secondary items can be declared as unknown when a simple mapping is in use.

The set of facilities available for data transfer control is itself one of the attributes of the model, so that where the more sophisticated features are not justified, they need not be implemented, and their absence can be made clear in the initial exchanges.

These simplifications mean that the bulk of the work in producing a minimal file transfer implementation over an existing transport station is likely to be the provision of local interfaces to the user and to his operating system.

In the long term, universal interworking must rest on the establishment of an internationaly agreed set of standards, and on the use of these standards by computer manufacturers. To bring the File Transfer Protocol to the attention of as

large an audience as possible, the document has been circulated widely, and presented before the INWG (IFIP 6.1). Moreover, there has recently been much activity within the standards organisations concerned with open working between computer systems; the relevant BSI committee will examine the protocol in the course of its deliberations and make its views known to the ISO. Such bodies will, doubtless, take due account of experience in real working systems, when comparing this solution with more tentative theoretical proposals.

P.F.Linington

COMMUNICATIONS HARDWARE

R.A.F. CHISHOLM, ERCC

H A R D W A R E

The last few years have produced major advances in data communications hardware and systems. Three areas of development in which the Communications Systems Division at Edinburgh Regional Computing Centre have been active are (a) the use of integrated circuits in digital data communications (b) the development of local private modems and line switching and patching facilities and (c) the development of a communications line monitor.

Communications Chips

1. UART - The advent of Large Scale Integrated (LSI) circuit technology heralded the introduction of digital communications orientated circuits. Around 1970 the bulk of computer digital data communications was concerned (and still is) with the transmission and reception of asynchronous data at fairly modest rates; it was therefore this area that the integrated circuit manufacturers first attacked. They produced a circuit known as a Universal Asynchronous Receiver/Transmitter (UART) which would handle the reception and transmission of serial line data and present a parallel data interface to the communications equipment. A typical UART is characterised in Fig 1. The first device to appear on the commercial market was a member of the General Instruments AY-5-1013 family. As Table 1 shows this device type has become the industry standard and is available from many manufacturers.

2. USRT - Around 1972 the first universal Synchronous Receiver/Transmitter (USRT) integrated circuit appeared. Advances in technology had produced a more complex chip which would run from a single power rail. A typical device is characterised in Fig 2. It does not appear that a true industry standard has emerged, as evidenced by Table 1.

3. USART - With further technological innovations the next step, reached around 1974, was the production of an integrated circuit which would in one package combine the functions of both the UART and the USRT. This chip known as the Universal Synchronous/Asynchronous Receiver/Transmitter (USART) is characterised in Fig 3. A innovation which appeared with these devices was the incorporation of modem control and status reporting features. Although there are a number of variants

of this device the Intel 8251 looks like being adopted as the industry standard.

4. MULTI-PROTOCOL - The acceptance of Synchronous Data Link Control (SDLC) communications procedures around 1974 provided the integrated circuit manufacturers with their next challenge in the data communications arena. The link level control procedures were ideally suited to implementation on a digital integrated circuit (Those of us who had to implement it in standard TTL parts can vouch for that !). The most significant offering in this field was released by Zilog in the third quarter of 1977. This device known as the Serial Input/Output (SIO) is characterised in Fig 4(a) and 4(b). Another quantum leap in technology has enabled Zilog to offer two full duplex channels, each capable of handling UART, USRT and HDLC type transmission on a forty pin chip. This must be vertical integration !

5. MODEM CHIP SETS - There are now available from various manufacturers, in the form of integrated circuits, many of the functional components which make up a modem. This is a fairly new area of development but indications from the USA suggest that many terminals are now being produced with integral modems. Fig 5 shows the six chip digital modem available from Cermetek.

For the future Zilog in announcing the Z8 microcomputer on a chip have heralded surely the next development - namely the communications controller as an integral part of the microcomputer chip. Fig 6 shows the proposed configuration of the Zilog Z8.

Switches and other network components

1. LOCAL USE PRIVATE MODEMS - As is often the case a technological advance in one area spawns, advances in other related fields - this is true of data communications chips. It is now relatively cheap to provide and buy data communications interfaces for most medium and small computer systems and hence there now exists a requirements to provide cheap "modems". At ERCC we have developed a range of asynchronous and synchronous baseband transceivers (modems) for use over limited distance private circuits. These units are considerably cheaper than commercially available narrow bandwidth

modems. Fig 7 characterises the synchronous units we have developed.

2. LINE PATCHING AND SWITCHING UNITS - In any large centralised data communications facility there exists a need to provide a degree of flexibility in the allocation of modems to terminals/computer ports. Being presented with these problems and unable to afford the high cost of proprietary items we developed a range of units which provide us with these features Fig 8 summarises the range of units, Figs 9(a) & 9(b) show the physical presentation of one of the types of unit.

Communications Line Monitor

During the development of communications hardware (Wideband Data Port) for Modular One computers, which act as nodes in the RCO network, it became obvious that it would be very useful if we could monitor both the transmitted and received data in a communications link between a computer port or terminal and a modem. Fig 10 illustrates the concept. A 'T'-piece is inserted in the receive and transmit data paths between the modem and the data terminal equipment. The 'T' piece which is totally passive, provides a mechanism for sampling both the data and control signals between the DCE and the DTE. The two information sources are fed into receivers in the monitor hardware. Fig 11 shows the various signals that are sampled on the receive and transmit paths. The monitor computer configuration is characterised in Fig 12. By far the major part of the project was the design and implementation of the software to synchronise, filter and analyse the information gathered from the lines; much effort went into the presentation of this information in a concise and meaningful manner. Space does not permit a lengthier description of the detail of this work but further information can be obtained from my colleague Ken Dietz who is responsible for the implementation of the project. Fig 13 characterises the next incarnation of hardware which is based on a DEC LSI-11 and is intended to provide a 'portable' system.

During my conversations over the two days there emerged a few points which are worth recording. A number of people have experienced problems in implementing communications chips in their designs, since

there appears to be some inconsistencies in manufacturers specifications, inabilities to perform to maximum published data rates, etc. I would be willing to act as a central collector of this information which I would pass on for publication in the Network News. Some of you may be aware that the American standard RS-232C has recently been superseded by RS-449 which presents a standard which in a fairly loose sense can be seen as presenting an interface with both balanced (CCITT V35) and unbalanced (CCITT V24) signalling capabilities.

Copies of this standard can be obtained from

EIA Engineering Department
Standards Orders
2001 Eye Street N.W.
Washington D.C. 20006
U.S.A.

The most recent price I have is $ 9:50

R.A.F. Chisholm
April 1978

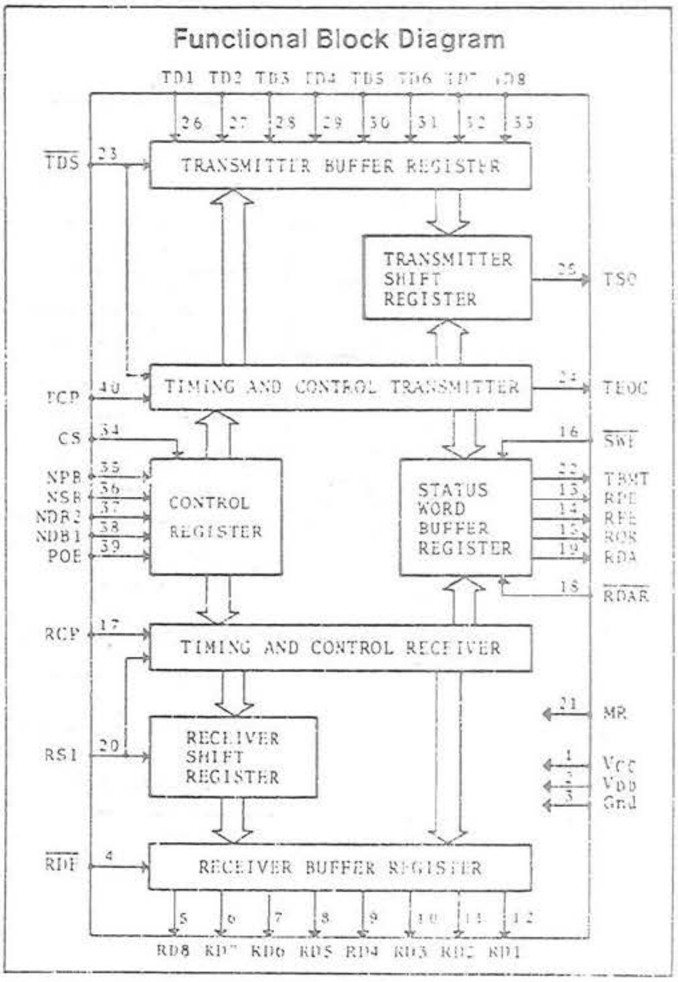

FEATURES

- **Direct TTL Compatibility**—no interfacing circuits required
- **Full or Half Duplex Operation**—can receive and transmit simultaneously at different baud rates
- **Fully Double Buffered**—eliminates need for precise external timing
- **Start Bit Verification**—decreases error rate
- **Fully Programmable**—data word length, parity mode, number of stop bits; one, one and one-half, or two
- **High Speed Operation**—40K baud, 200ns strobes
- **Master Reset**—Resets all status outputs
- **Tri-State Outputs**—bus structure oriented
- **Low Power**—minimum power requirements
- **Input Protected**—eliminates handling problems
- **Hermetic Dip Package**—easy board insertion

UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER (UART)

<u>Figure 1</u>

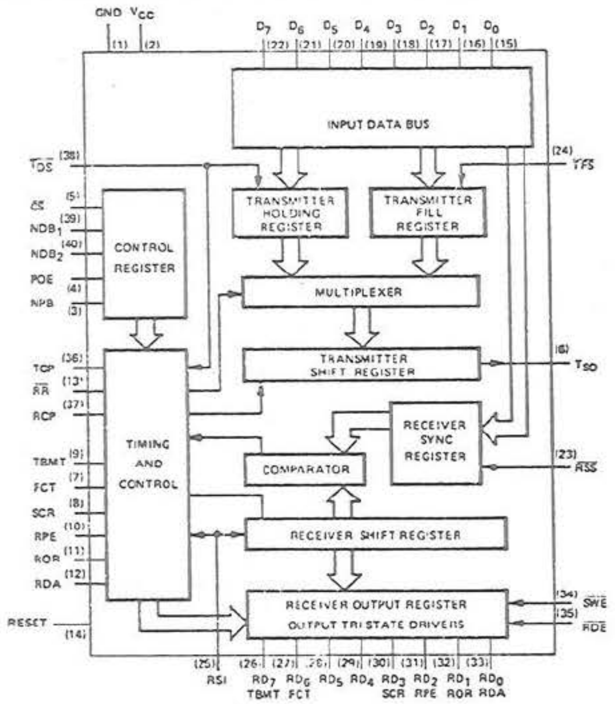

BLOCK DIAGRAM

FEATURES

- 500 KHz Data Rates
- Internal Sync Detection
- Fill Character Register
- Double Buffered Input/Output
- Bus Oriented Outputs
- 5–8 Bit Characters
- Odd/Even or No Parity
- Error Status Flags
- Single Power Supply (+5v)
- Input/Output TTL Compatible

UNIVERSAL SYNCHRONOUS RECEIVER/TRANSMITTER (USRT)

Figure 2

MANF - TYPE	DATA RATE	POWER	PINS	FEATURES	PRICE
UART					
GI AY-5-13	0-30	+5-12	40		£3.83
TI TMS6011	0-12	+5-12	40	Equivalent to above	£4.96
AMI S1883	0-12	+5-12	40	" " "	-
WD TR1402	0-30	+5-12	40	" " "	£6.45
WD TR1602	0-30	+5-12	40	" " "	£3.40
SMC COM2017	0-40	+5-12	40	" " "	
SMC COM2502	0-40	+5-12	40	" " "	
WD TR1863	0-40	+5-12	40	" " "	£4.96
WD TR1983	0-9.6	+5	28	MDM Control & Status	
MOT MC6850	0-.5M.	+5	24	MDM Control & clock 1x16x64	£7.20
AMI S6850				As Above	
USRT					
SMC COM2601	0-250K	+5-15	40		
AMI S 2350	0-500K	+5	40		
MOT. MC6852	0-600K	+5	28	MDM Cont. & Status	10.41
ASTRO					
WD UC1671	0-1MHZ	+5+12-5	40	Clock Rates - Add Decode	£29.34
WD UC1971	56 SYN 9.6 ASYNC	"	"	MDM Control	£24
WD TR1953	"	+5	28	MDM Control Clk 1 x 16 x 64	
INTEL C8521	"	+5	28	As above & Test	£7.17
SIG 2651	0-.8M	+5	28	16 Clock Rates (ASYNC)	£13.06
MULTI-PROTOCOL					
WD SD1933	0-1.5M	+5	40		£40.33
SMC COM5025	0-2M	+5	40		£60
SIG 2652	0-.5M	+5	40		£43.61
MOT. MC6854	0-.5M	+5	40		15.56
Z160G S10	o-.5M	+5	40		£60

COMMUNICATIONS CHIPS

TABLE 1

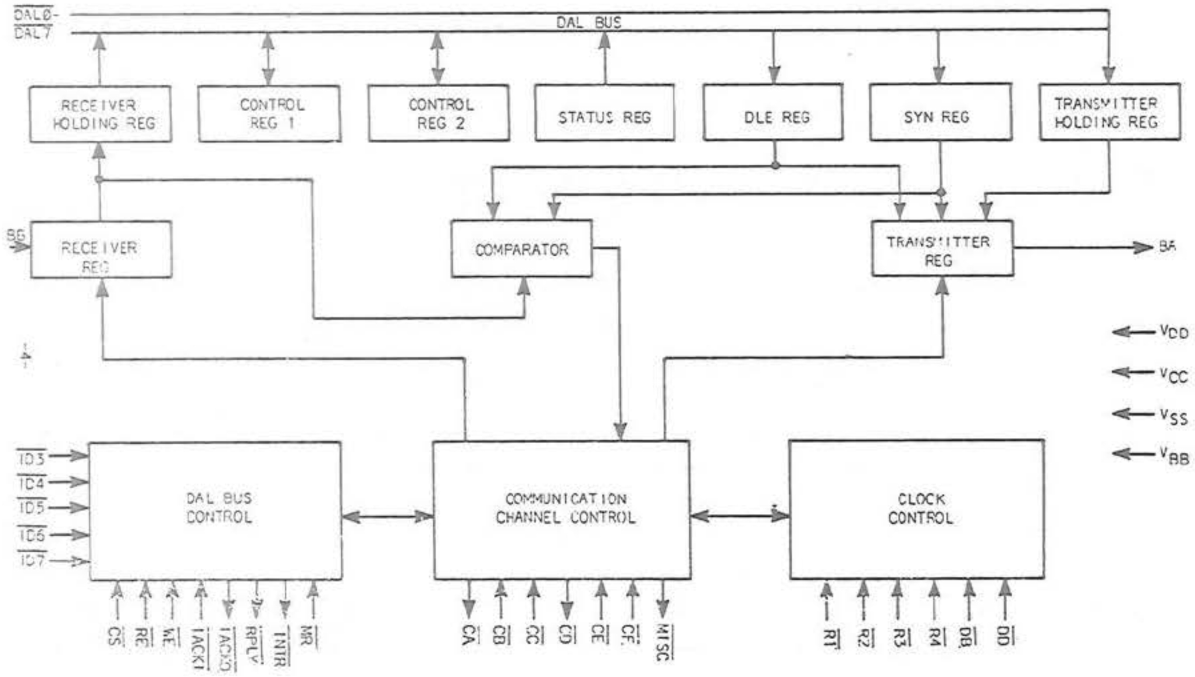

FEATURES

- SYNCHRONOUS AND ASYNCHRONOUS
  Full Duplex Operations
- SYNCHRONOUS MODE
  Selectable 5-8 Bit Characters
  Two Successive SYN Characters Sets Synchronization
  Programmable SYN and DLE Character Stripping
  Programmable SYN and DLE-SYN Fill
- ASYNCHRONOUS MODE
  Selectable 5-8 Bit Characters
  Line Break Detection and Generation
  1-, 1½-, or 2-Stop Bit Selection
  False Start Bit Detection
  Automatic Serial Echo Mode
- BAUD RATE - DC TO 1M BAUD/SEC
- 8 SELECTABLE CLOCK RATES
  Accepts 1X Clock and Up To 4 Different 32X Baud Rate Clock Inputs
  Up To 47% Distortion Allowance With 32X Clock
- SYSTEM COMPATIBILITY
  Double Buffering of Data
  8-Bit Bi-Directional Bus For Data, Status, and Control Words
  All Inputs and Outputs TTL Compatible
  Up To 32 ASTROS Can Be Addressed On Bus
  On-Line Diagnostic Capability

TRANSMISSION ERROR DETECTION-PARITY
Overrun and Framing

UNIVERSAL SYNCHRONOUS/ASYNCHRONOUS/RECEIVER TRANSMITTER (USART)

Figure 3

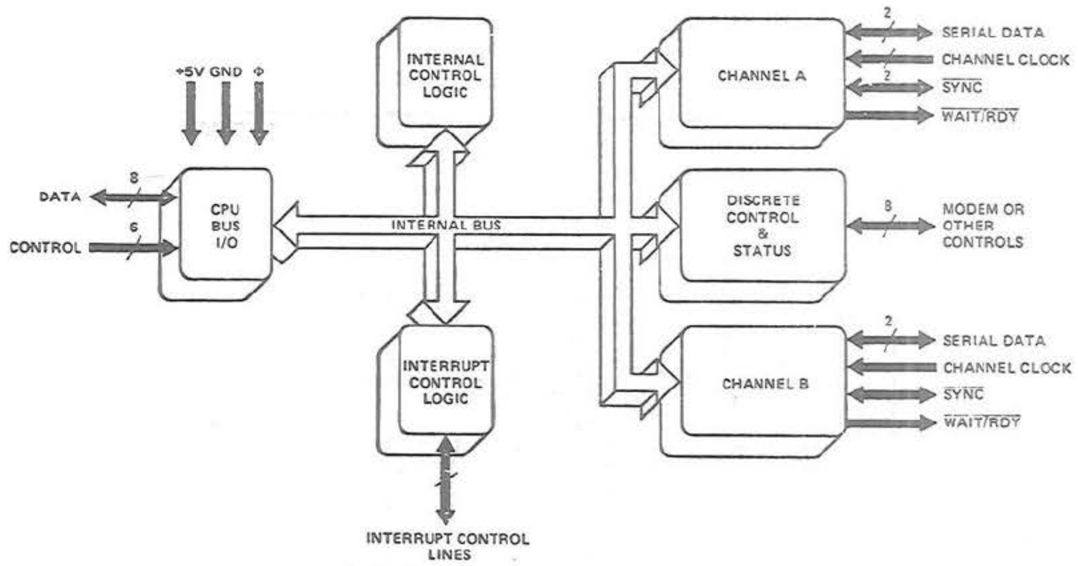

FIGURE 1
SIO BLOCK DIAGRAM

Features

- Two independent full duplex channels
- Data rates – 0 to 550K bits/second
- Receiver data registers quadruply buffered; transmitter doubly buffered.
- Asynchronous operation
  - 5, 6, 7 or 8 bits/character
  - 1, 1½ or 2 stop bits
  - Even, odd or no parity
  - x1, x16, x32 and x 64 clock modes
  - Break generation and detection
  - Parity, Overrun and Framing error detection
- Binary Synchronous operation
  - Internal or external character synchronozation
  - One or two Sync characters in separate registers
  - Automatic Sync Character Insertion
  - CRC generation and checking.
- HDLC or IBM SDLC operation
  - Automatic Zero insertion and deletion
  - Automatic Flag insertion
  - Address field recognition
  - I-field residue handling
  - Valid receive messages protected from overrun
  - CRC generation and checking
- Eight modem control inputs and outputs
- Both CRC-16 and CRC-CCITT (-0 and -1) are implemented
- Daisy chain priority interrupt logic included to provide for automatic interrupt vectoring without external logic.
- All inputs and outputs fully TTL compatible.

MULTI-PROTOCOL CONTROLLER - ZILOG SIO

Figure 4 (a)

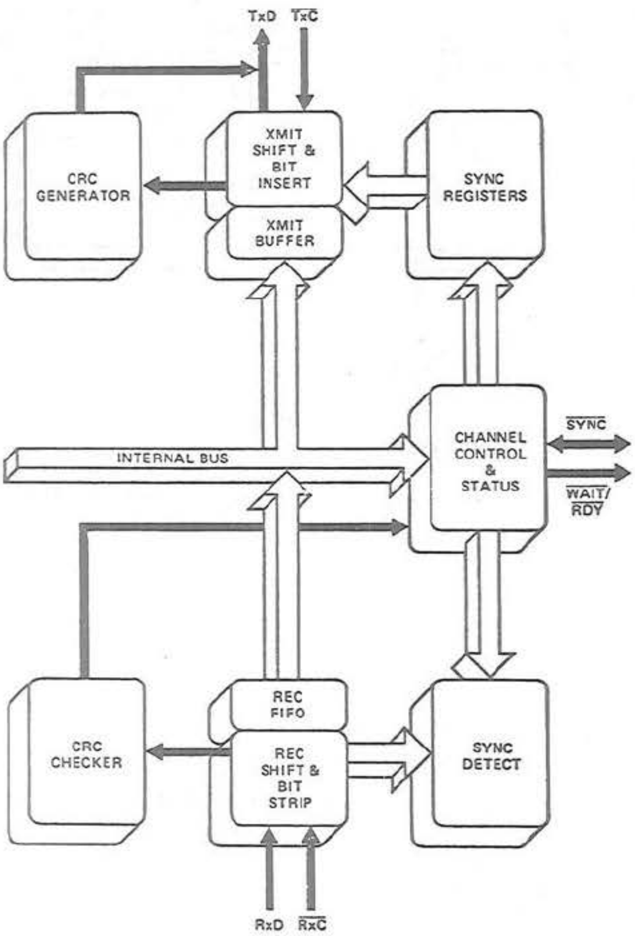

FIGURE 2
CHANNEL BLOCK DIAGRAM

MULTI PROTOCOL CONTROLLER - ZILOG SIO

Figure 4 (b)

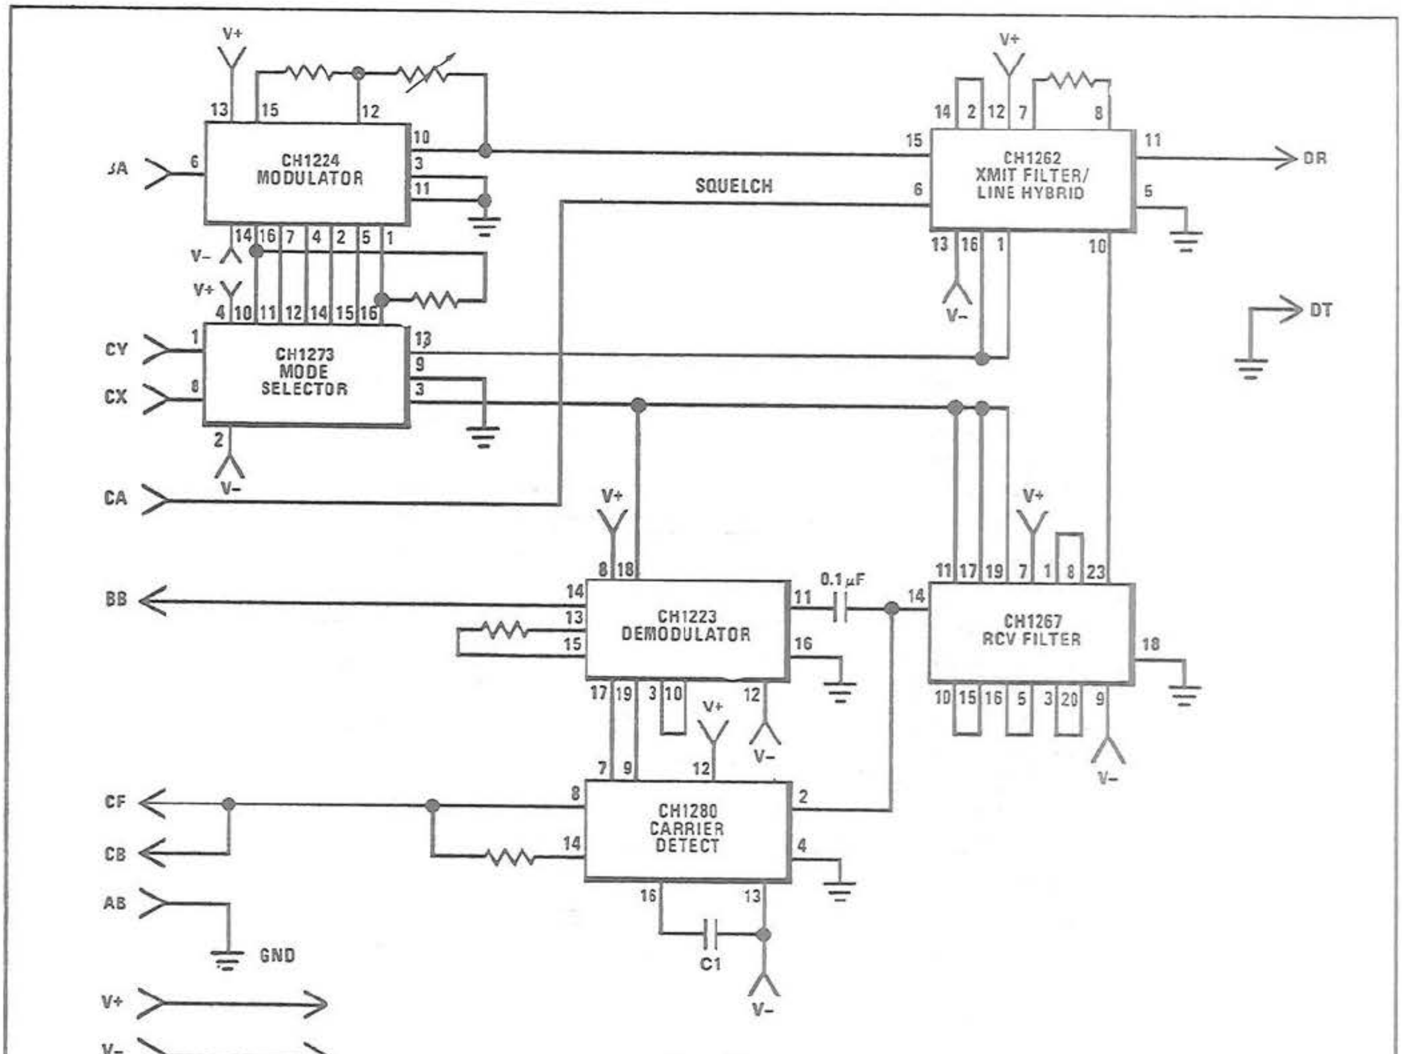

Figure 1

Data Format	0 to 300 BPS, Serial, Binary Asynchronous
Operation	Full duplex over switched telephone network
Modulation	Phase-coherent FSK
Frequency Tolerance	±0.5% max. (standard) ±0.2% max. (optional)
Transmitter Output Impedance	600 Ω
Transmitter Output Level	0dBm to –12dBm
Receiver Dynamic Range	–50dBm to 0dBm
Bit Error Rate	< 1 x 10^{-5} for 8dB s/n –50dBm receive level 3002 unconditioned line, –8dBm adjacent channel
P-P Jitter	5% max.
Receive Clamp	Carrier detect "OFF" clamps RECEIVED DATA in "MARK"
Carrier Detect Threshold	"ON" at –51dBm, "OFF" at –54dBm
Carrier Detect Timing	"OFF" to "ON" – 150 ± 50ms "ON" to "OFF" – 50 ± 25ms
Data and Control Interfaces	RS-232-C
Max. Supply Voltage	±18 V_{DC}
Power Consumption	1.4W typical
Output Short Circuit Duration	indefinite
Operating Temperature Range	0°C to +70°C
Storage Temperature Range	–25°C to +85°C

FREQUENCIES

		Originate	Answer
Transmit	"MARK"	1270Hz	2225Hz
	"SPACE"	1070Hz	2025Hz
Receive	"MARK"	2225Hz	1270Hz
	"SPACE"	2025Hz	1070Hz

MODEM CHIP SET

Figure 5

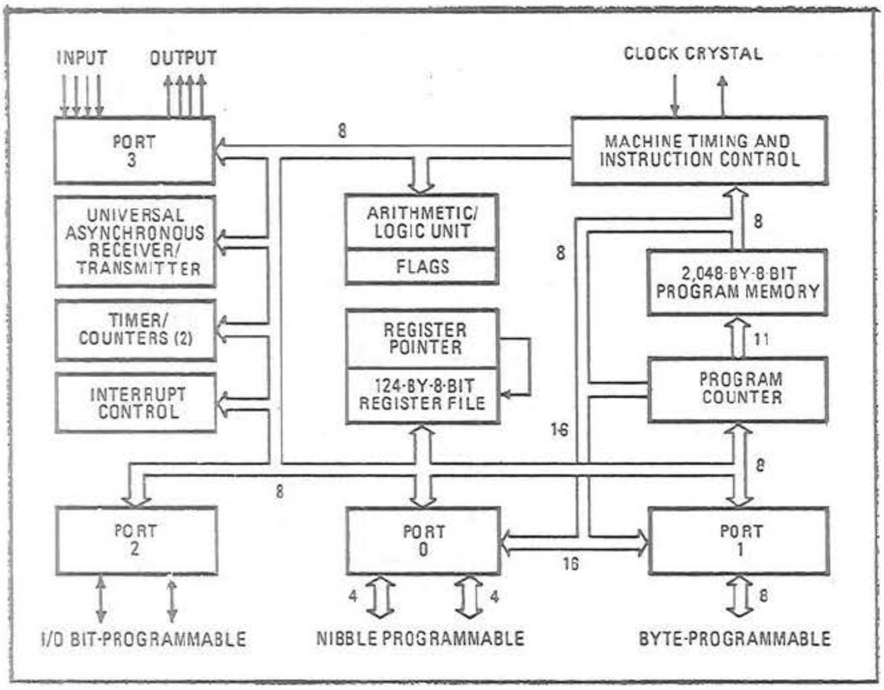

ZILOG Z8 MICRO COMPUTER - BLOCK DIAGRAM

<u>Figure 6</u>

SYNCHRONOUS BASEBAND TRANSCEIVERS

FEATURES:

* LOW COST
* OPERATES OVER 4-WIRE PRIVATE CIRCUITS
* DI-PHASE MODULATION
* DIFFERENTIAL LINE DRIVING
* FOUR SWITCH-SELECTABLE SPEEDS
    a. 2.4 Kb/s
    b. 4.8 kb/s
    c. 9.6 Kb/s
    d. 19.2 Kb/s
* SELECTABLE TIMING SOURCE
    a. Master ('A' end) - Internal
    b. Slave ('B' end) - i. External
    ii. 'A' end sourced & reconstituted
* MODULAR DESIGN
    a. Based on Eurocard format
    b. Available in two formats
        i. 4 channel box
        ii. 16 channel crate
* FOUR INDICATORS
    a. Tx Data
    b. Rx Data
    c. CTS
    d. DCD
* LOW POWER CONSUMPTION - CMOS based
* USES PLL FOR ACCURATE TIMING RECOVERY
* OPERATING RANGE (o.5mm Solid Wire)
    a. 6.0km @ 2.4 Kb/s
    b. 4.8km @ 4.8 Kb/s
    c. 3.5km @ 9.6 Kb/s
    d. 3.0km @ 19.2 Kb/s

Figure 7

UNIVERSITY OF EDINBURGH

Edinburgh Regional Computing Centre

Communications Interface Flexibility Modules

This display features a range of modules which have been designed to provide the patching and switching features often required in a digital data communications environment. The modules are designed to cater for both CCITTV24 and V35 data communications interfaces and are built to Type 62 equipment standards.

The range of modules produced include

- **Patching Module Type A**

  CCITT V24 - 12 circuit patching

  Normal through jacks connect modem to computer port

  Any computer port may be patched to any modem.

  8 circuits per module

- **Patching Module Type C**

  CCITT V24 - 3 circuit patching (RX,TX common)

  Normal through jack connects modem to computer port

  Computer ports connection via multi-way connector

  Any computer port may be patched to any modem

  8 circuits per module

- **Switching Module Type B**

  CCITT V24 - 12 circuit switching

  Modem can be switched to one of 3 computer ports

  4 modem circuits per module

- **Switching Module Type E**

  CCITT V24 - 12 circuit switching

  Computer port can be switched to one of 3 modems

  4 computer ports per module

EDINBURGH REGIONAL COMPUTING CENTRE
JAMES CLERK MAXWELL BUILDING, THE KING'S BUILDINGS, MAYFIELD ROAD, EDINBURGH, EH9 3JZ. 031-667 1081

Figure 8

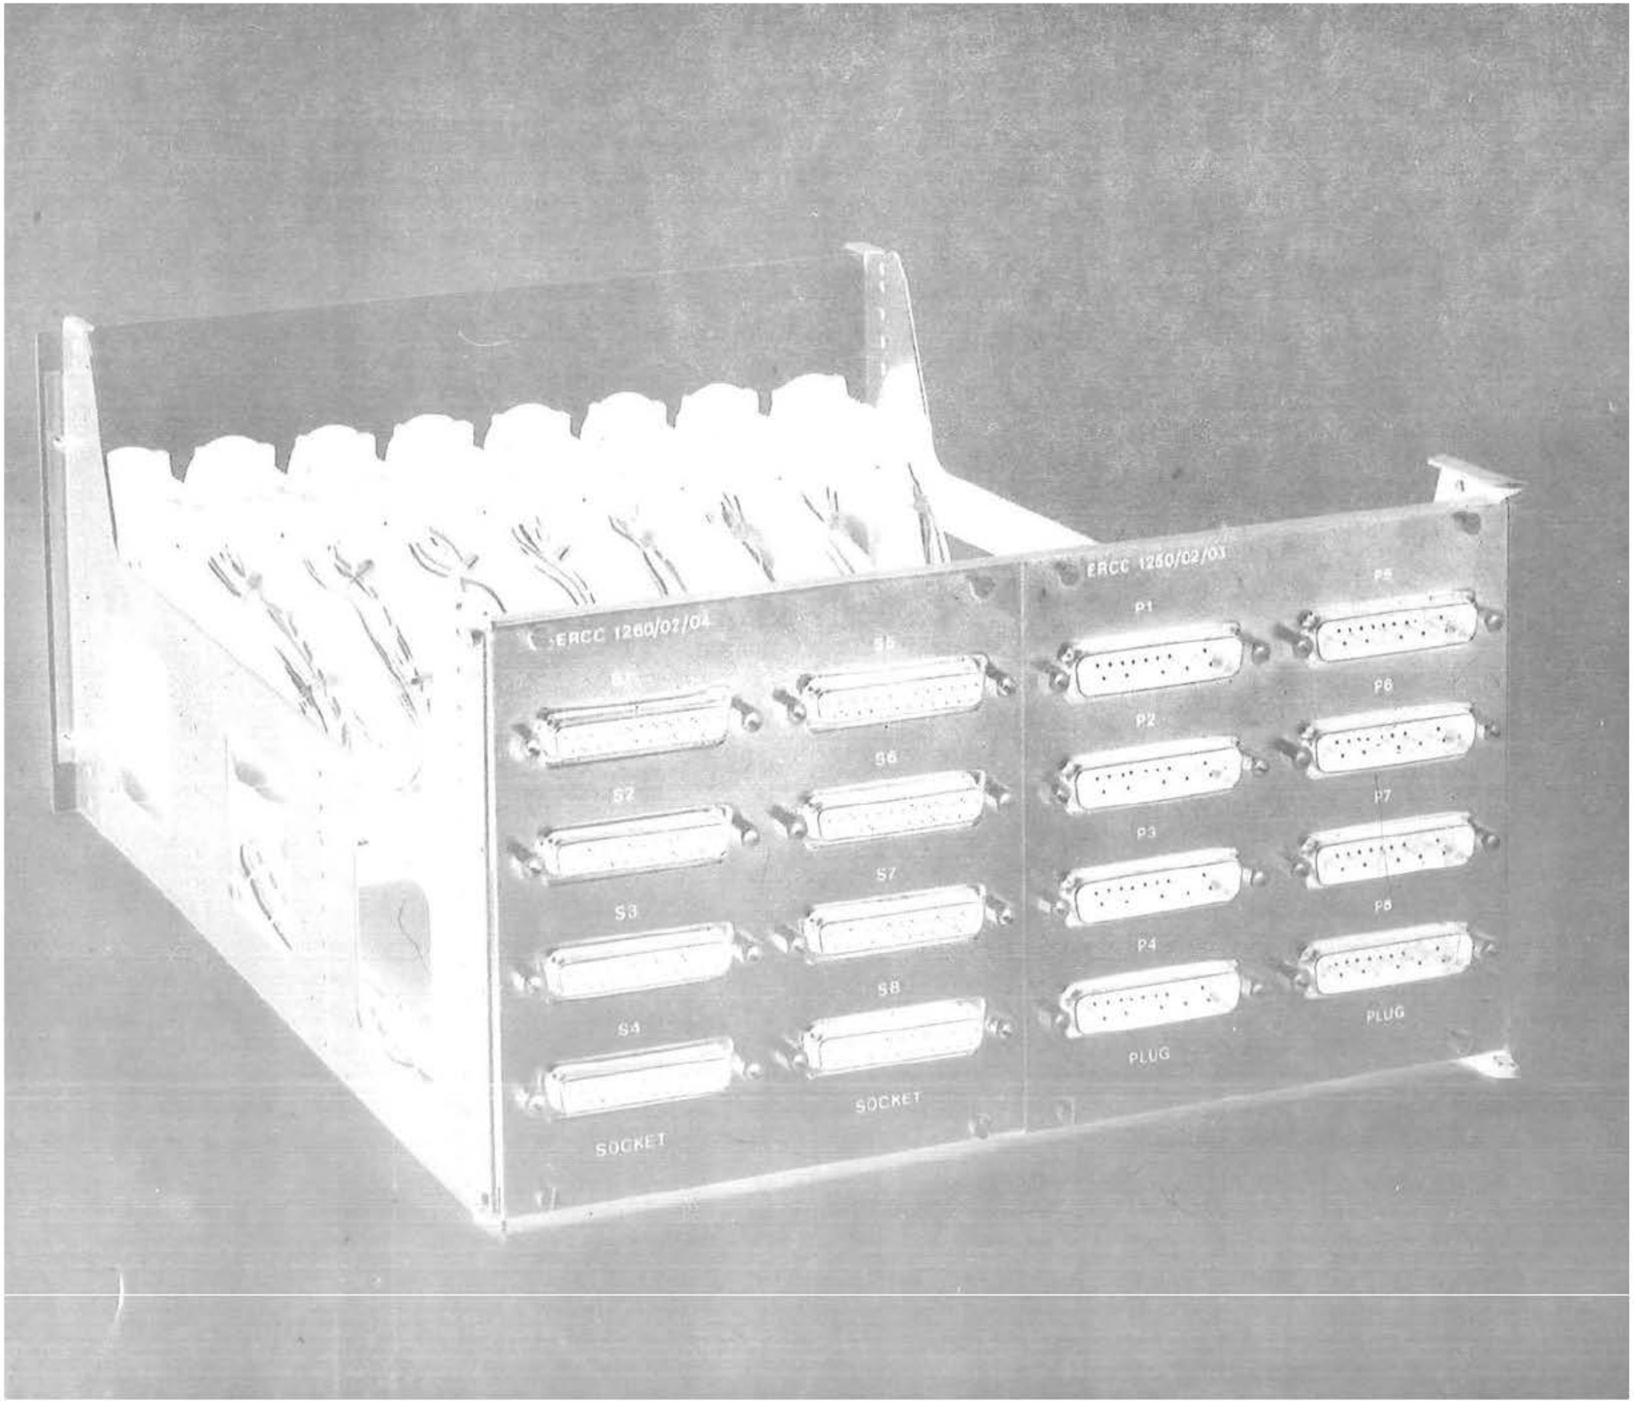
ERCC 1260/02/04
S5
S2
S6
S7
S3
S8
S4
SOCKET
SOCKET
ERCC 1260/02/03
P1
P5
P6
P2
P7
P3
P8
P4
PLUG
PLUG

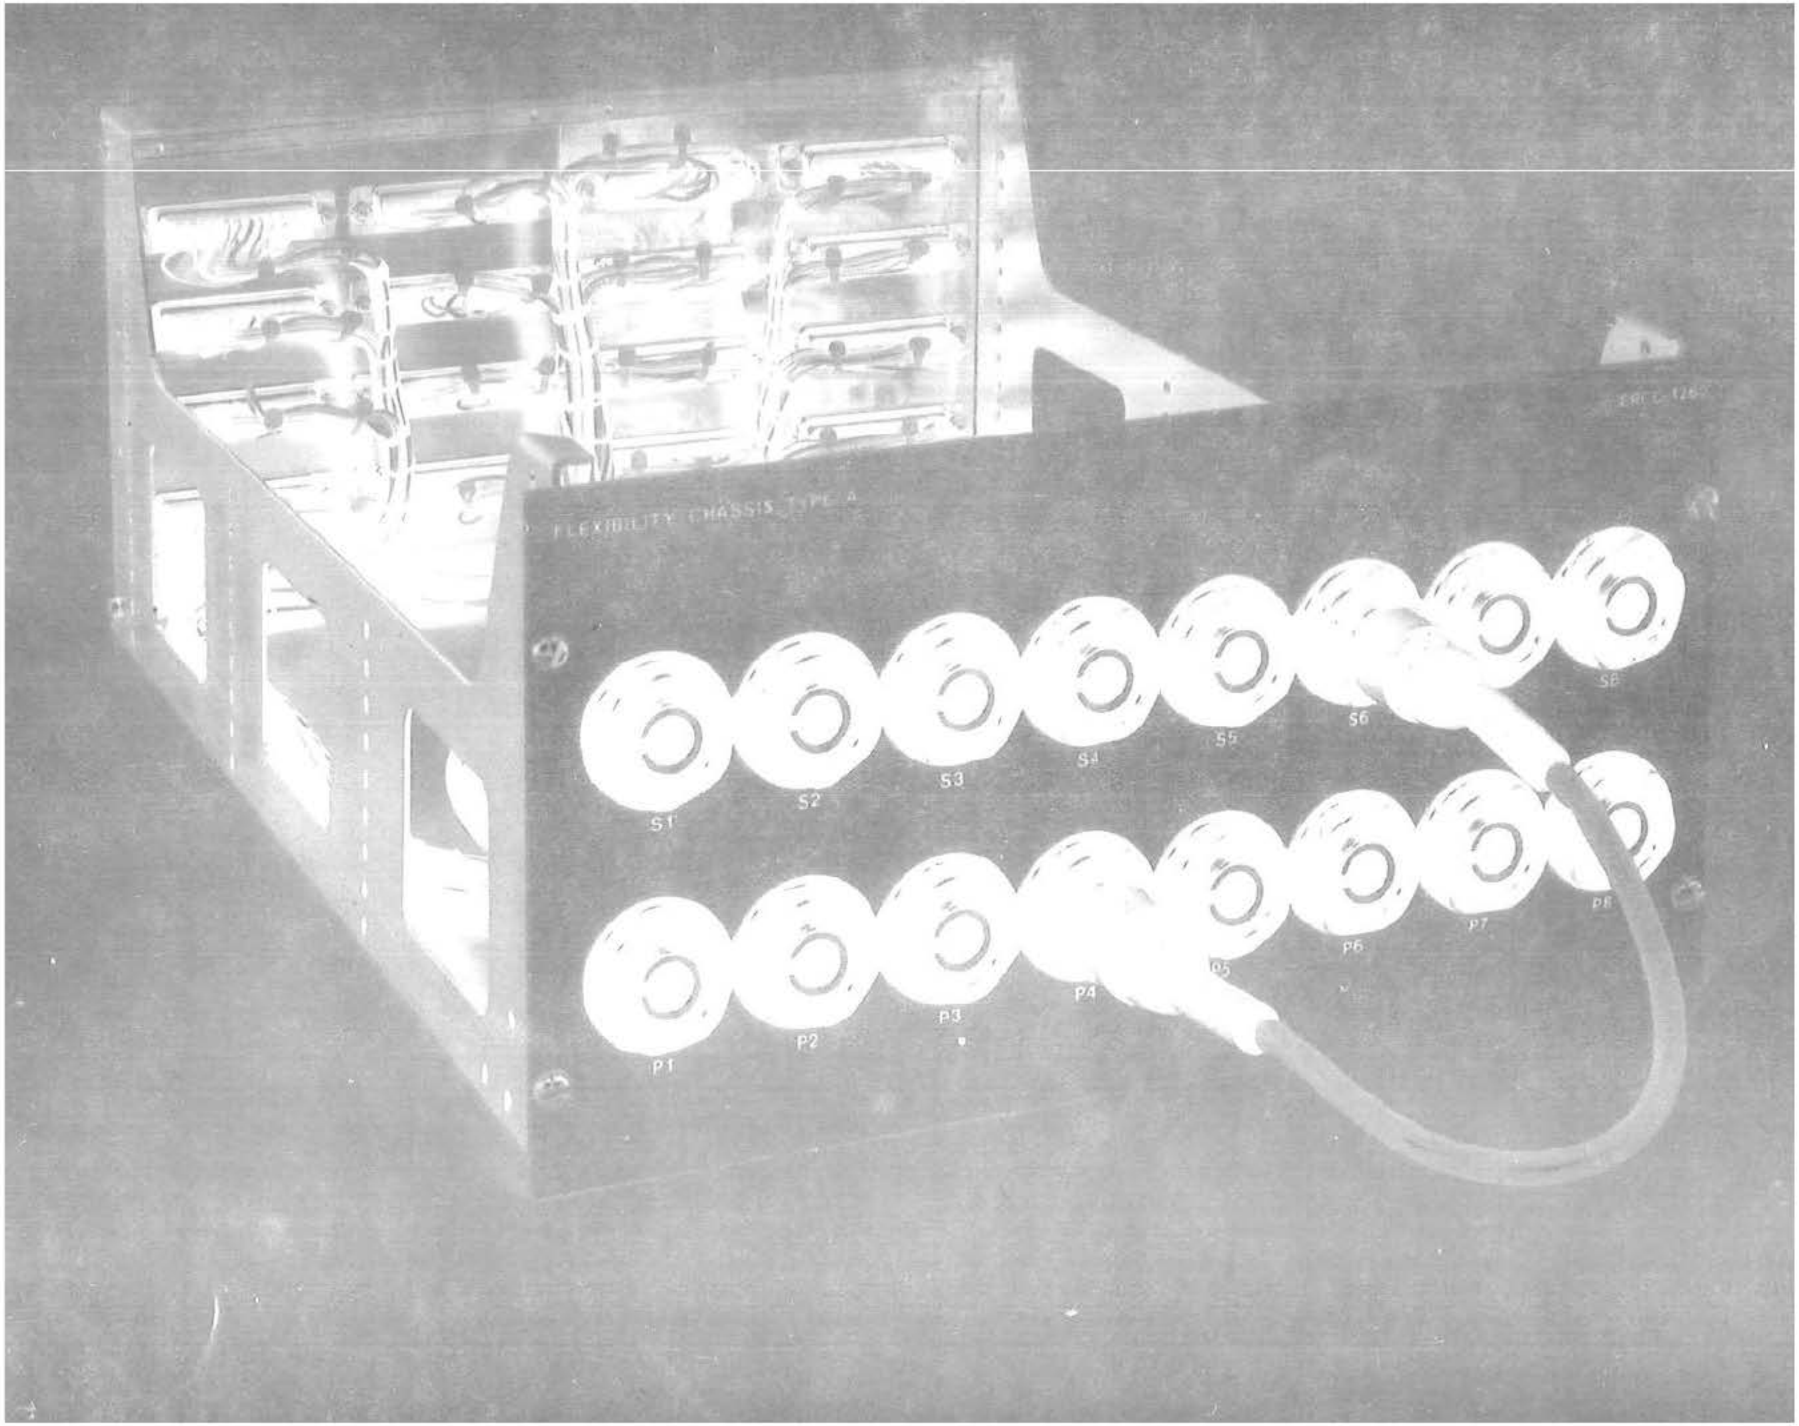
FLEXIBILITY CHASSIS TYPE A
S1
S2
S3
S4
S5
S6
S8
P1
P2
P3
P4
P5
P6
P7
P8

DATA TRANSMISSION

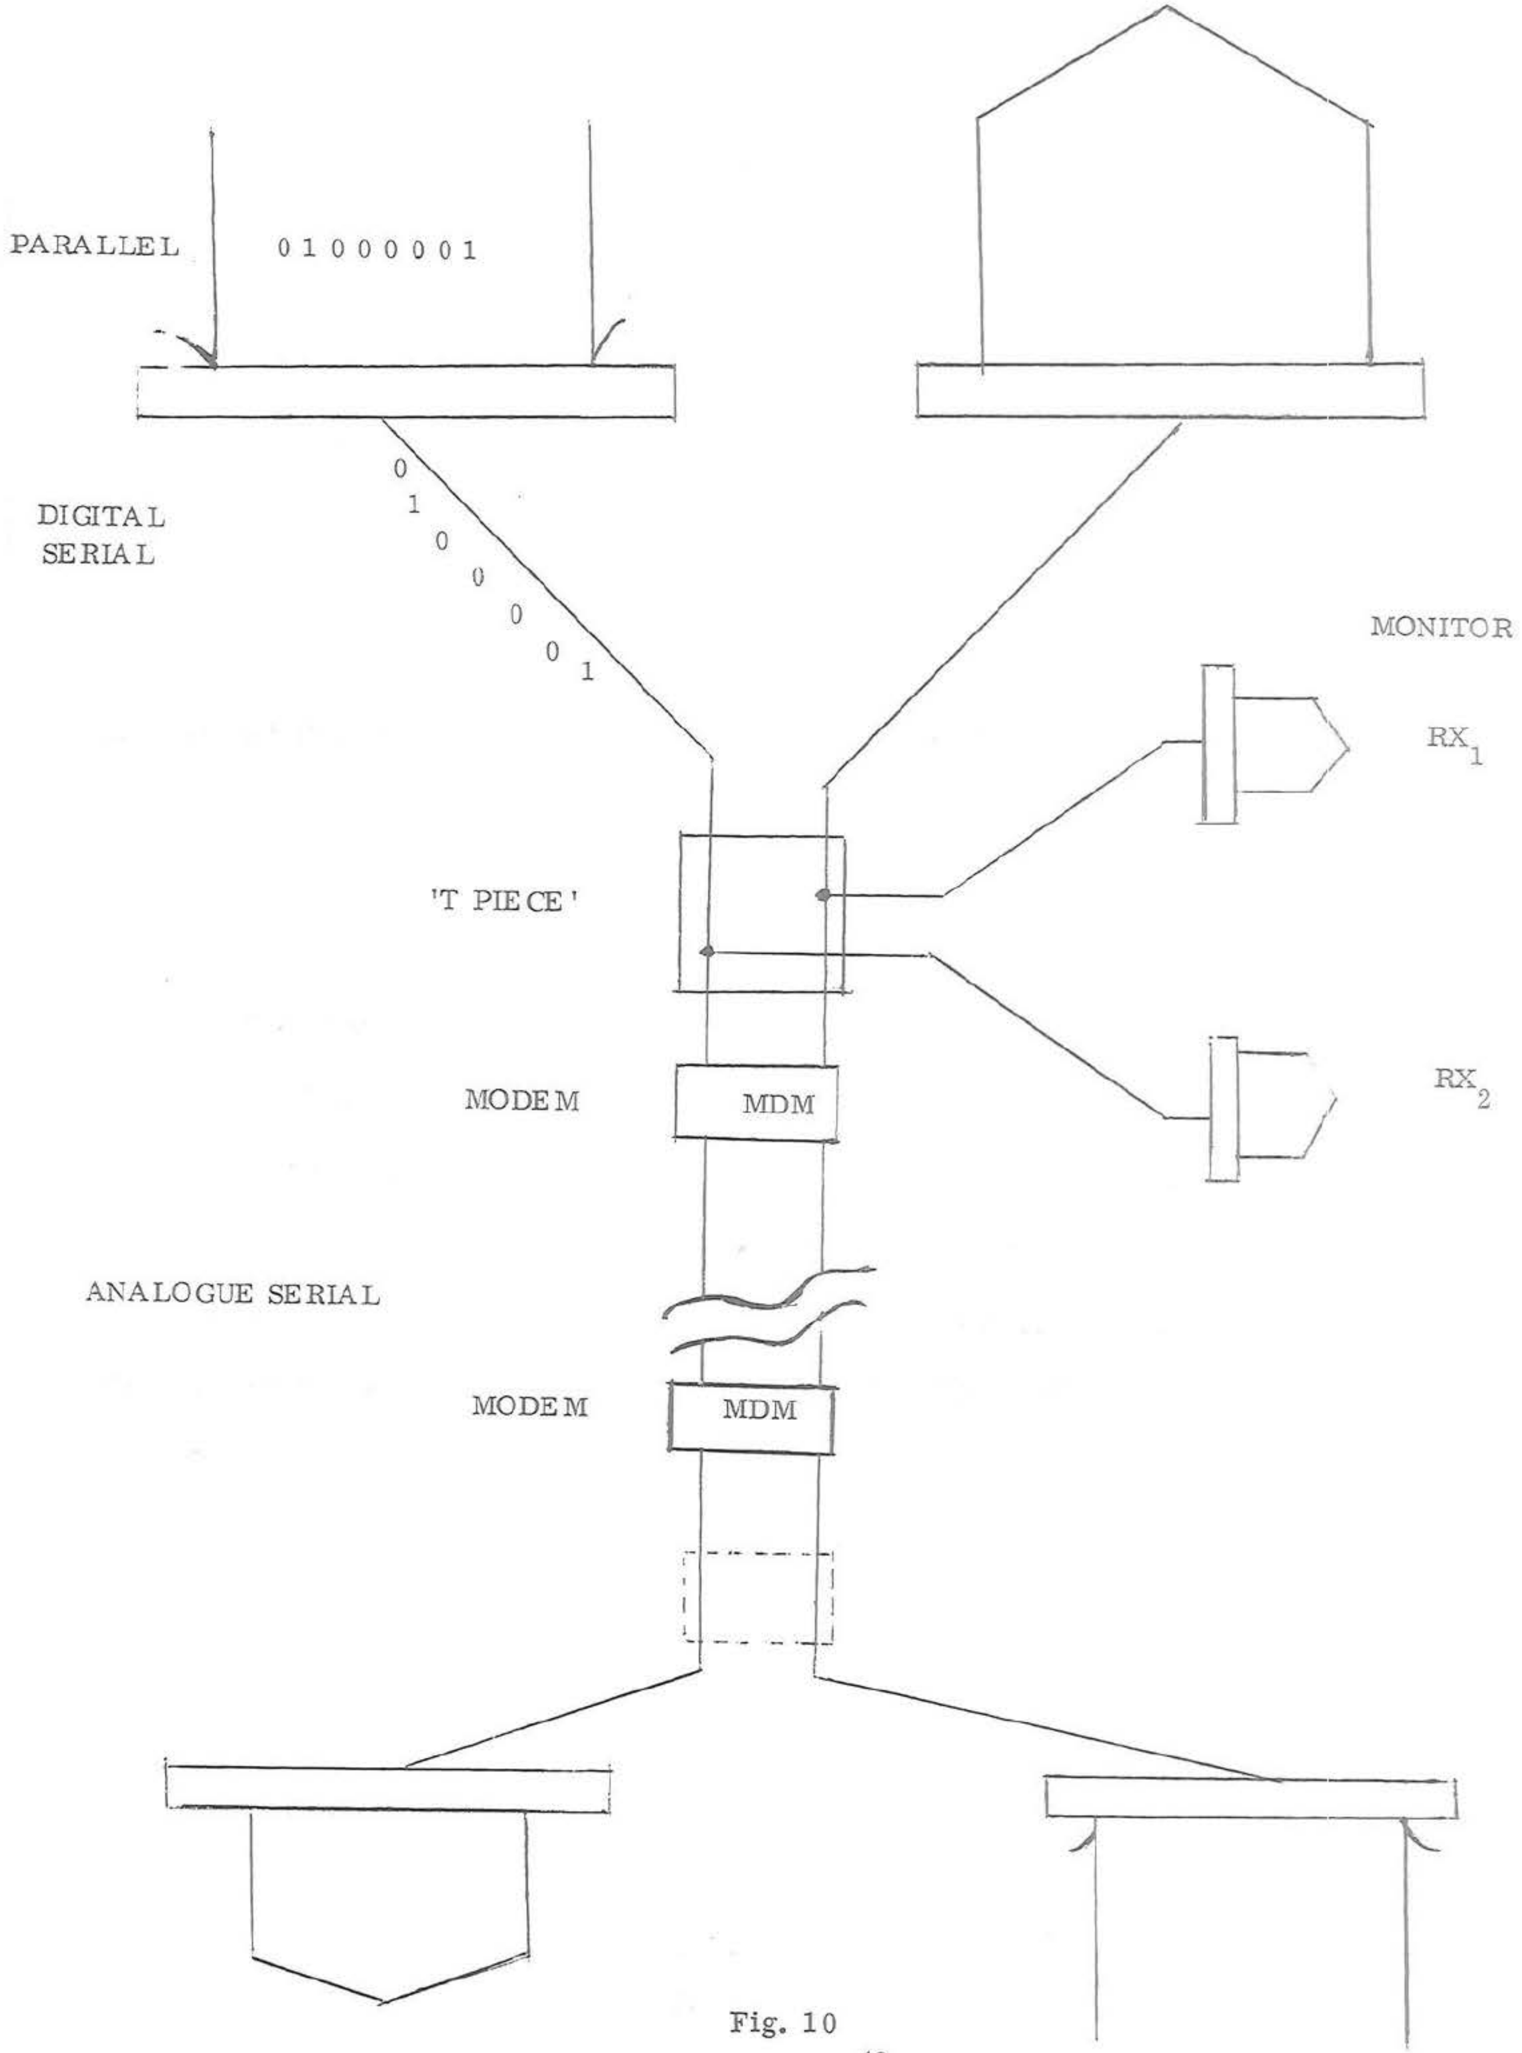

Fig. 10

SIGNALS COPIED

DTE	MONITOR Rx1
Rx DATA	•
Rx CLOCK	•
DATA SET READY	•
DATA CARRIER DETECT	•
(Rx CLOCK DETECT)	CALLING INDICATOR

DCE	MONITOR Rx2
Tx DATA	Rx DATA
Tx CLOCK	Rx CLOCK
DATA SET READY	
READY FOR SENDING	DATA CARRIER DETECTED
(Tx CLOCK DETECT)	CALLING INDICATOR

Fig. 11

MONITOR MACHINE HARDWARE COMPLEMENT

1 x DEC PDP 11-10 PROCESSOR

- 20 KWORDS CORE STORE (16 BIT)
- 110 BAUD CURRENT LOOP CONSOLE INTERFACE
- 20 MSEC CLOCK

1 x DEC RK11/05 DISC (EXCHANGEABLE CARTRIDGE)

- 1.2 MWORD CAPACITY

2 x SYNCHRONOUS INTERFACE (SINGLE CHANNEL)

- V24
- BYTE OPERATION ONLY
- ERCC DESIGN

2 x ASYNCHRONOUS INTERFACE (SINGLE CHANNEL)

- V24
- UP TO 96KBAUD
- DEC DL11E

1 x TPIECE (ERCC)

Fig. 12

UNDER DEVELOPMENT

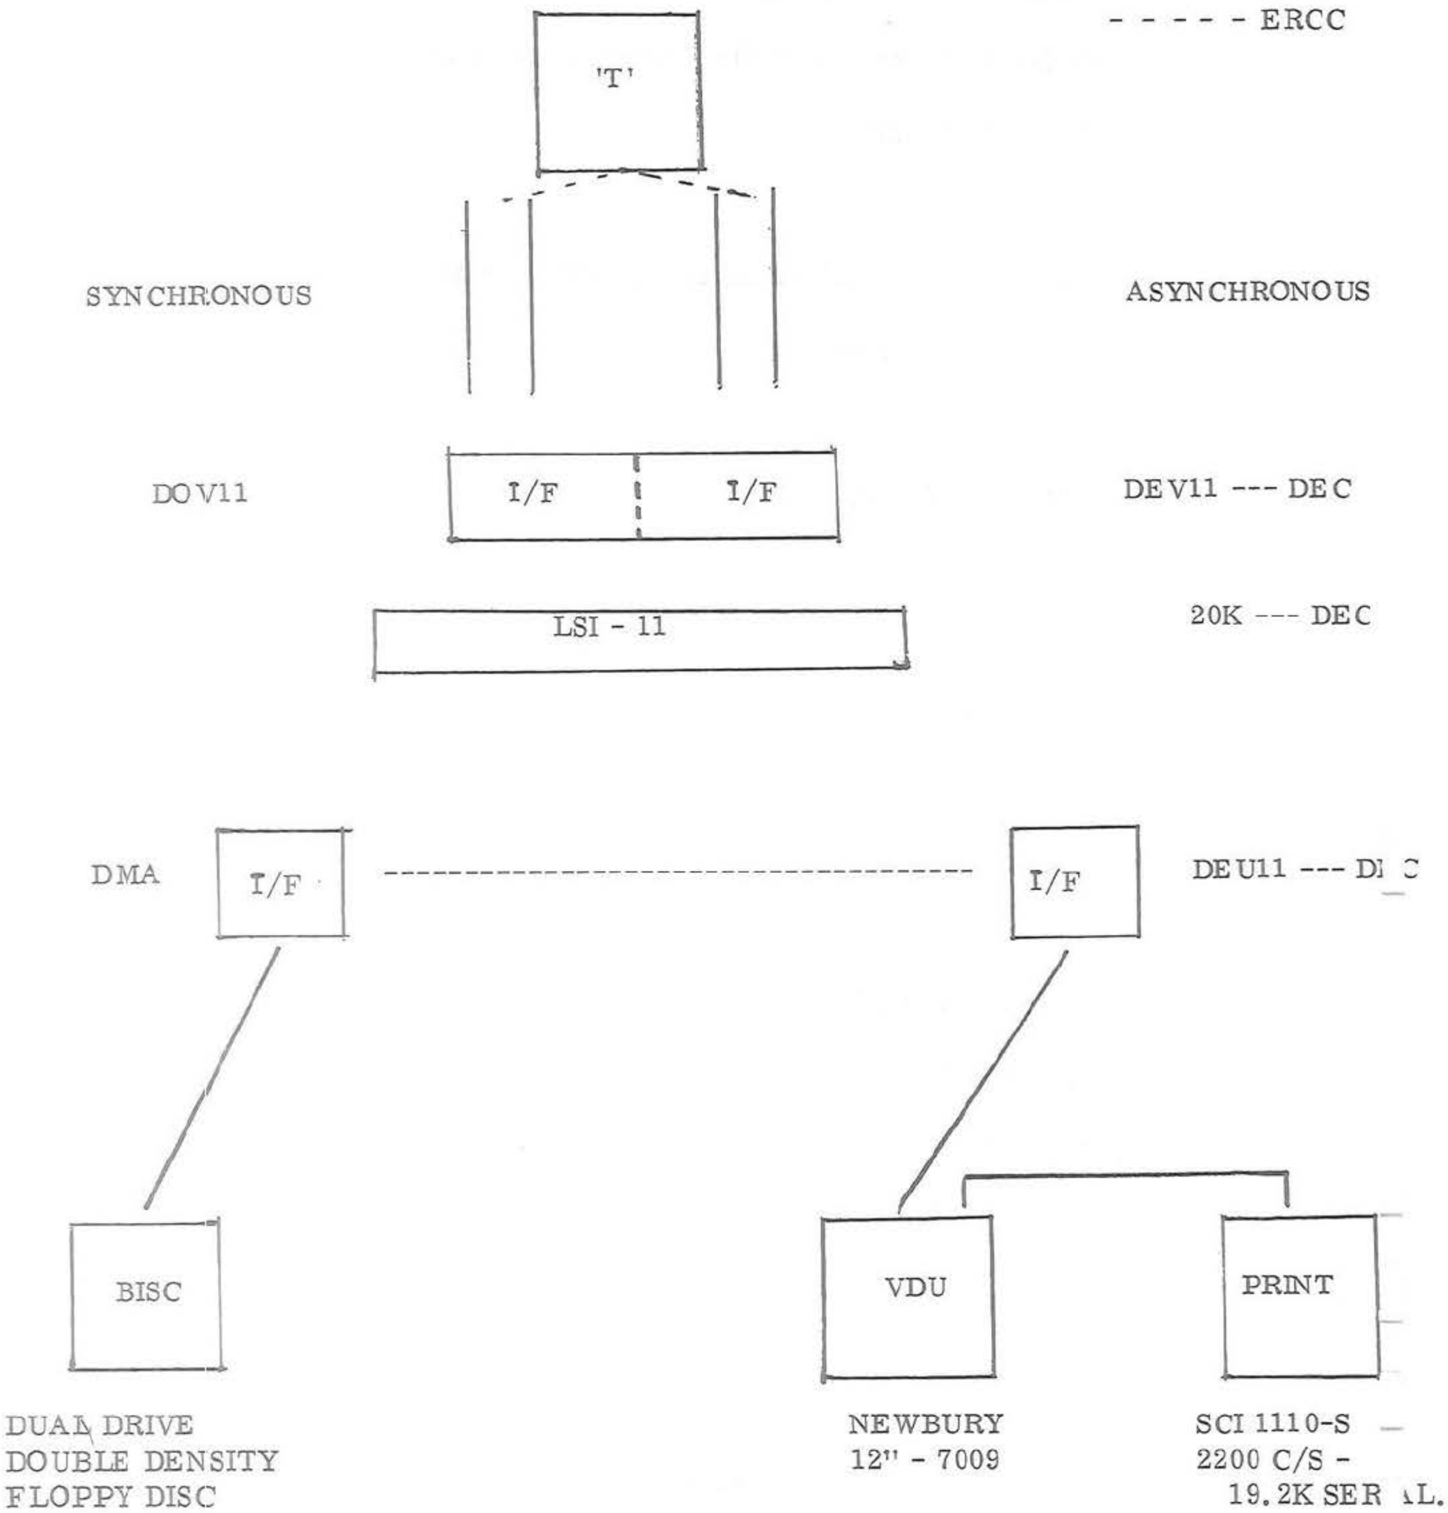

APPROX. COST - £11,000

Fig. 13

REGIONAL NETWORK ACTIVITIES

METRONET - London Region	J.P. Brandon, QMC
RCO	J.I. Davies, ERCC
South-East Universities	G. Litchfield, Oxford
MIDNET - Midlands Region	M.A. McConachie, Nottingham
SWUCN - South West Region	J. Thomas, SWURCC
SRC	A. Peatfield, SRC Daresbury
Report on the North West Network Service	J. Lindley, Salford

METRONET - London Region

Presenter: J. P. Brandon

Present configuration:

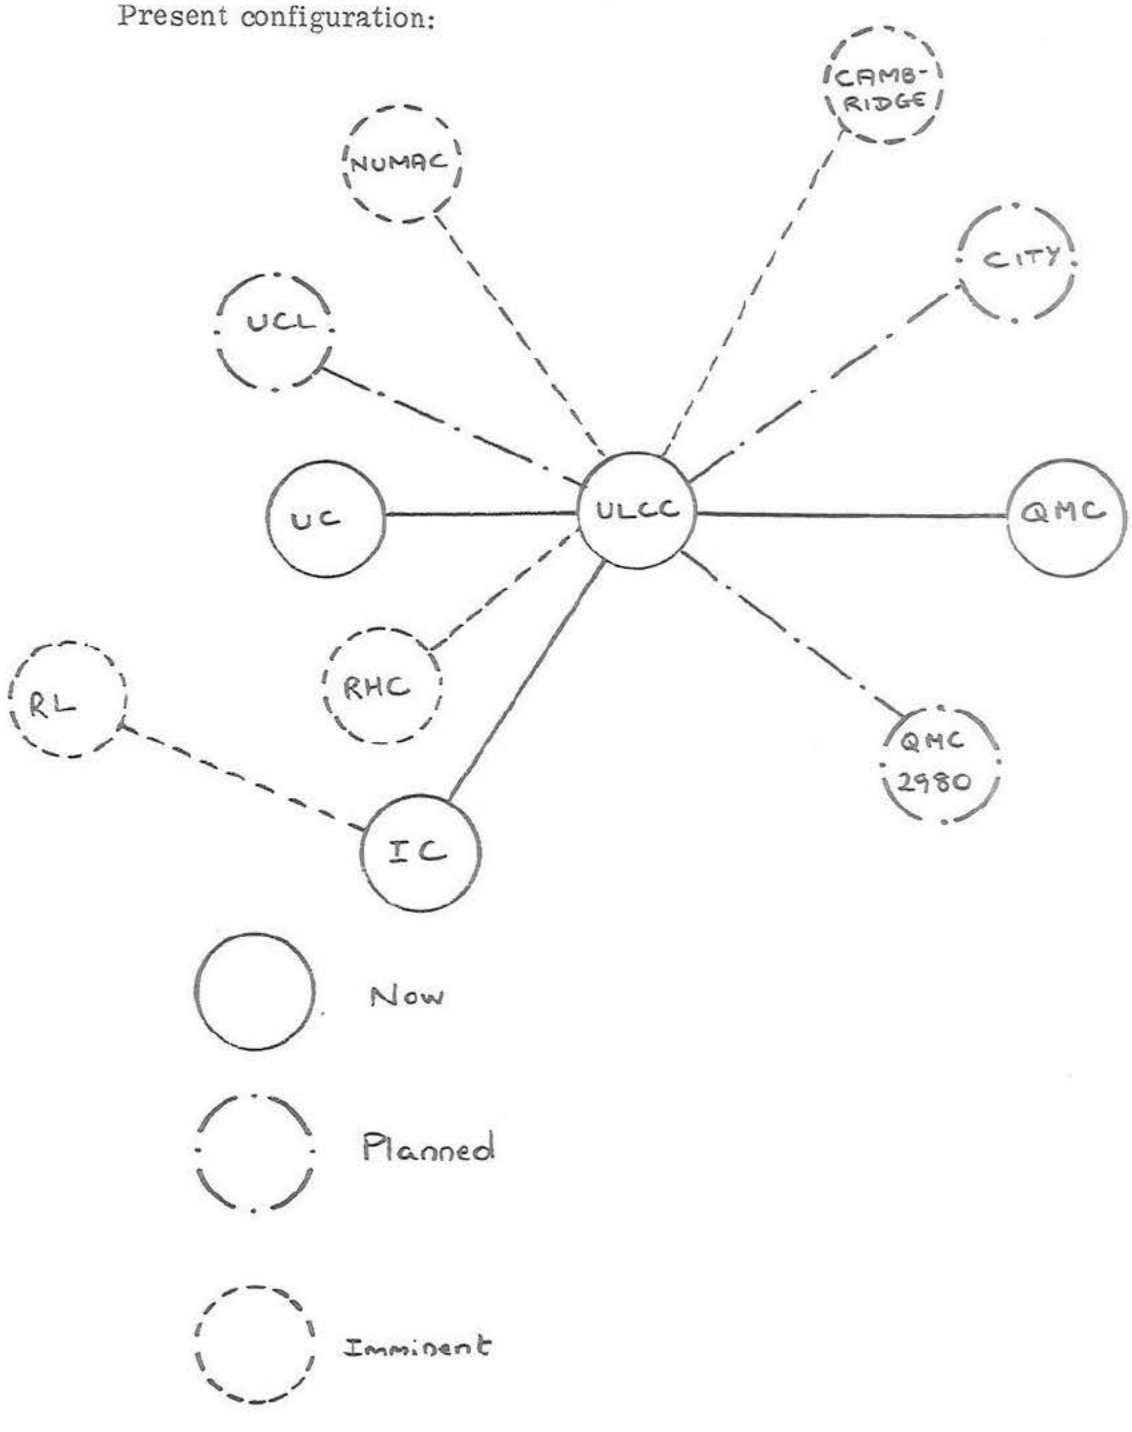

The traffic level is currently at several hundred jobs/day moving around the network. New documentation is being written:

a. an operator document explaining what different sites input/output looks like;
b. a common user document aimed at Advisory Services entitled "What is Metronet!"

Job tracing is a problem; it is difficult for operators and users to trace jobs in an ad-hoc network.

Currently METRONET II is being planned. It is hoped to solve network problems for the next 10 years. One of the major problems in the new design will be to maintain the existing services. Studies of expected traffic point to very high traffic levels. London have been talking to the P.O. with a view to the P.O. running an X25 PSE for them.

RCO

Presenter: J. Davies

Current Configuration

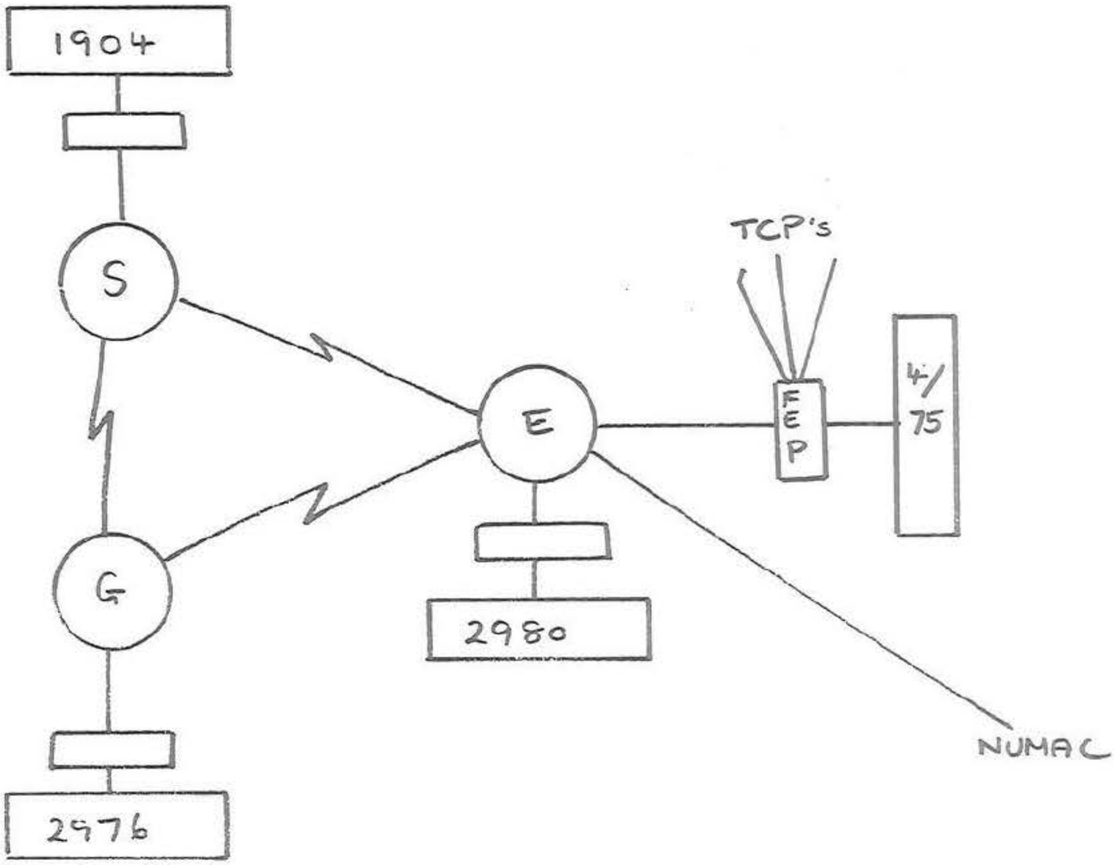

Strathclyde and Glasgow run during the day shift, and Edinburgh 24 hours a day.

Current Work

1. Hardware survey (current Mod. 1 nodes 7 years old)
2. TSS DCP for Glasgow and Bush estate 2980.
3. Exec rewrite and software restructure for second Edinburgh node about 2/3 having to be rewritten because the throughput is being affected by delays in sending packets through the switch.
4. ITP to and from NUMAC.
5. TCP connections to Network.

The developed configuration, expected in about one year, will be:

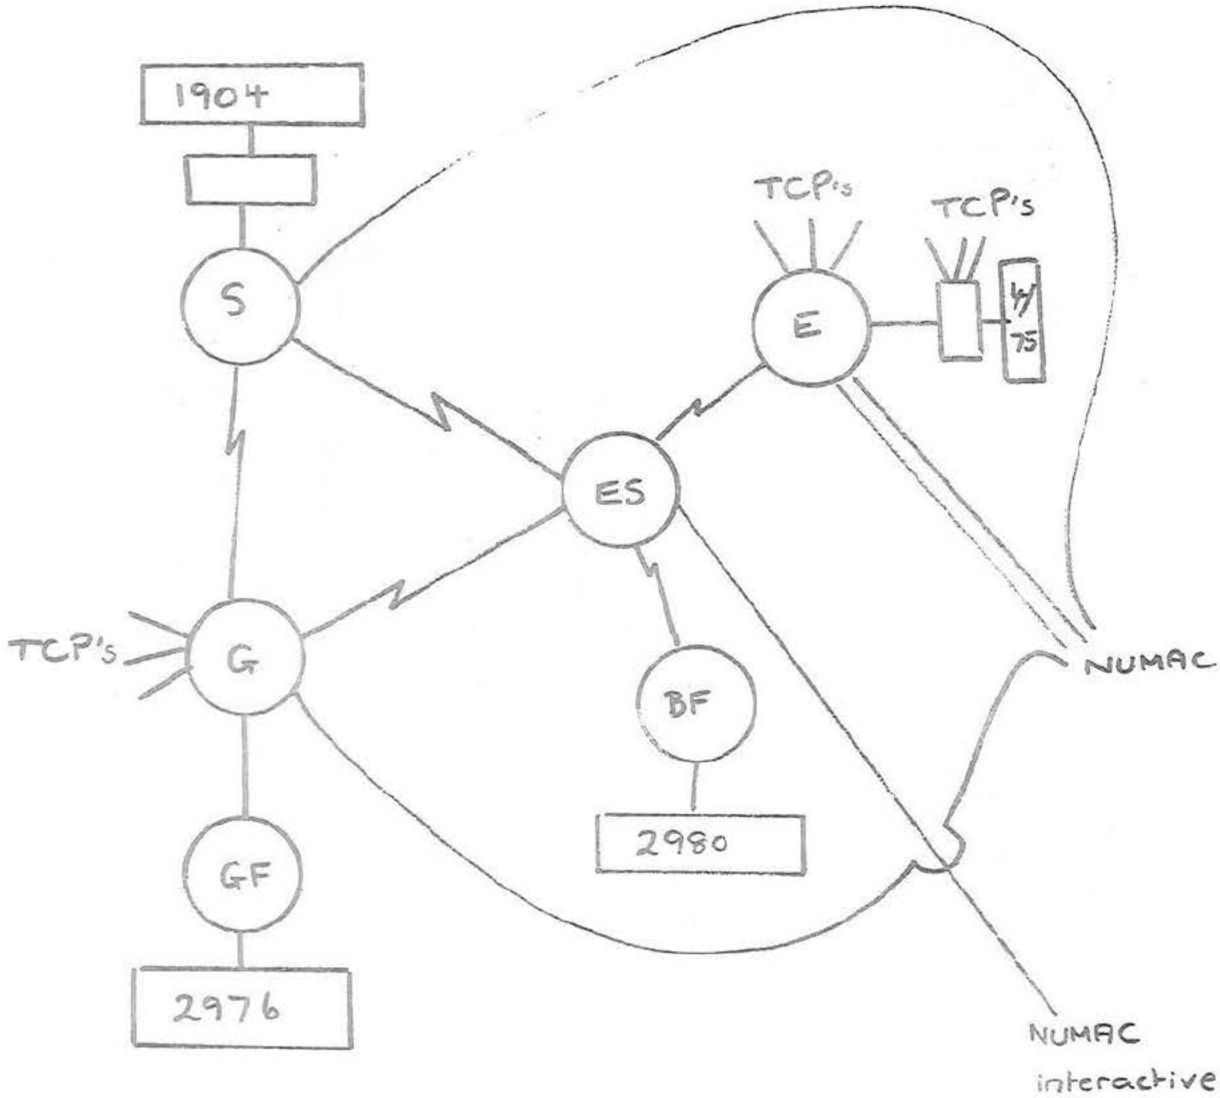

The RCONET currently has three Nodes in operation with Node-to-Node communication established between Strathclyde, Glasgow and Edinburgh. Links to the three local mainframes are established although the 2976 at Glasgow is not yet in service. Links to both regional mainframes, NUMAC and the ICL 2980 at Edinburgh, are established.

We are currently working on a number of enhancements to the network service. A survey of current switching, interfacing, RJE and terminal Concentrator machinery is in progress. This is not intended to parallel the Post Office provision of a public network or the Network Unit's attempts to obtain standard machines for switching and interfacing. As providers of an existing service, we feel we must make plans for its future based on our own knowledge of the directions of current technology.

We are continuing work on a software module to communicate directly with a 2900 VME/B mainframe (TSS DCP). We are attempting to improve the performance of existing Nodes by rewriting our executive and rationalising the module structure within the Nodes. We have agreed to connect our network to the NUMAC interactive network which is already using our NSI (link and call level) and ITP (high level interactive terminal) protocols. Our current interactive terminal concentrators (TCP) use these protocols and we are in the process of attaching them directly to the network rather than via the Edinburgh EMAS front end machine.

It is intended that the network shall expand to include three more Nodes two of these will also act as front-ends to the Glasgow 2976 and Edinburgh 2980. Components of a redundant Satellite One will form a new Node at Edinburgh. The Strathclyde 1904 FEP is to be modified to use NSI procedures for network communication.

S.E. Universities

Presenter: G. Litchfield

Grouping includes Southampton, Sussex, Kent, Surrey, Reading, Brunel, City, Oxford. Sites have links to ULCL, UMRCC, Newcastle, Cambridge and Rutherford. The 2980 at Oxford recently started a service of which 40% is available for the region with current links to Kent and Southampton.

S.E. Universities have a committee which is currently discussing the building of a packet-switched network based on the NPL packet switch specification. They might be in favour of having a P.O. run regional P.S.E. because of the amounts of work required to interface to an X25 network.

The objectives of the Network are:

a. package sharing,
b. resource sharing,
c. database access,
d. rationalisation of access to remote sites,
e. ready access to national facilities,
f. resilience.

MIDNET - Midlands Universities

Presenter: M.A. McConachie

Current links

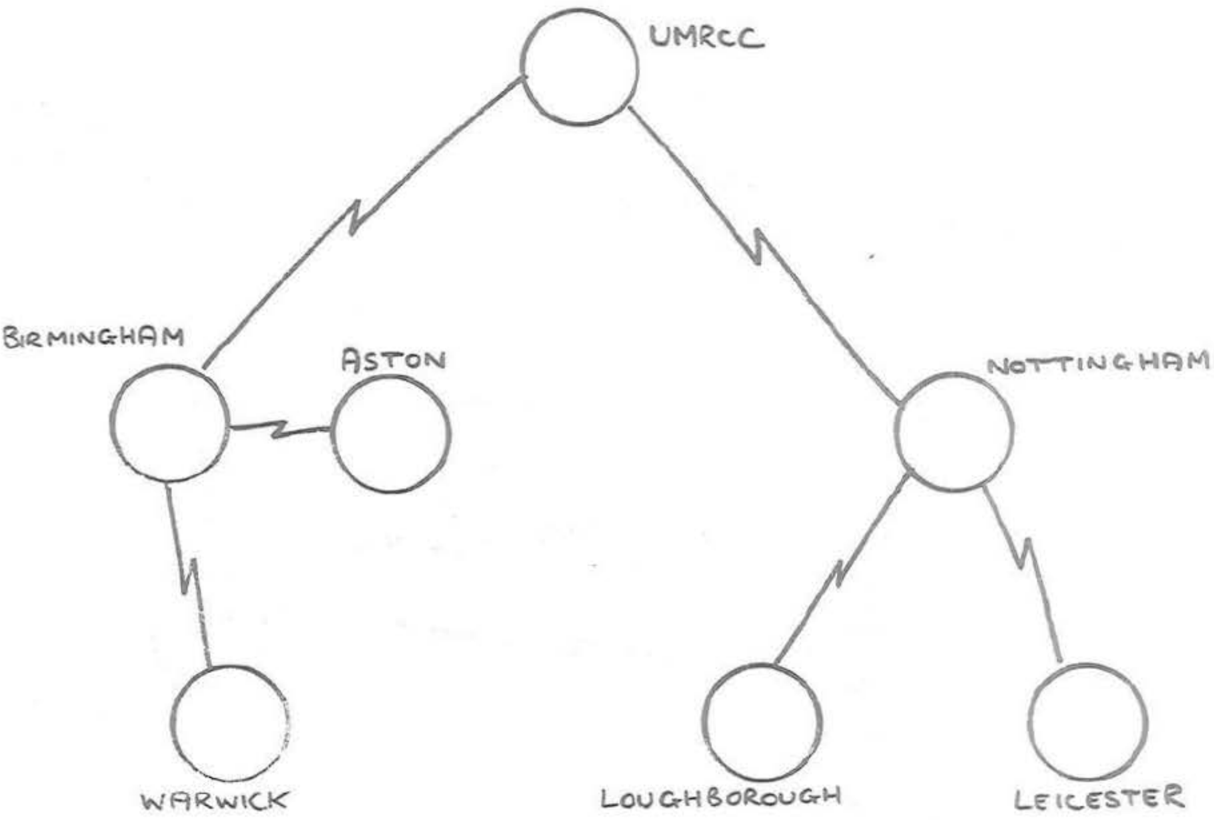

The network is used for access to UMRCC only with no connections between machines.

The same objectives for a future network as S.E., timescale for MIDNET is the next 3 or 4 years.

MIDNET Configuration

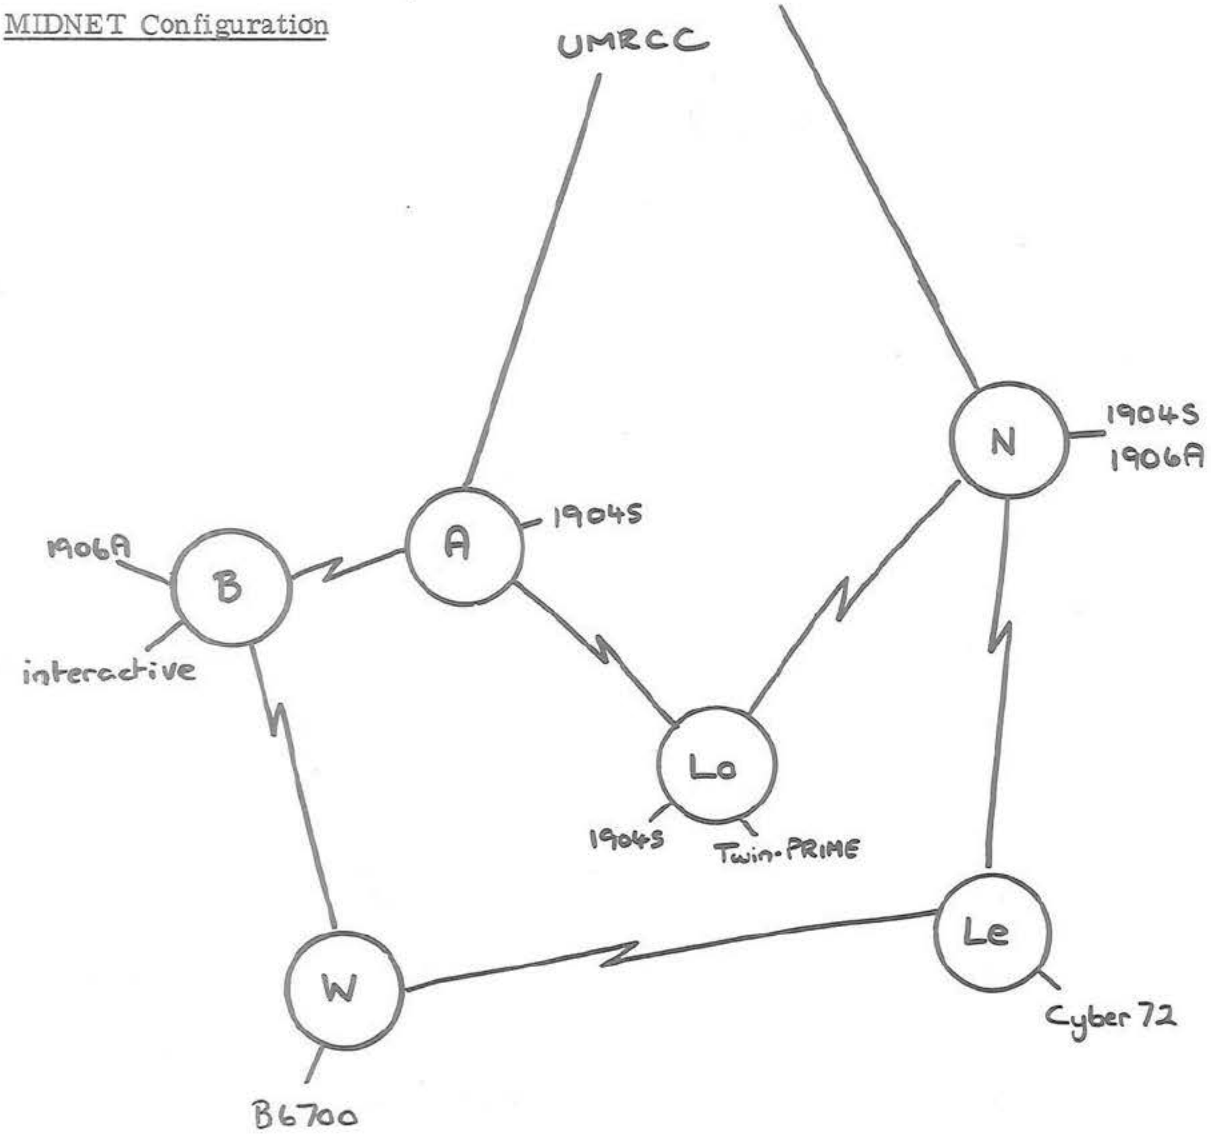

Nodes are PDP 11/34 supplied by SYSTIME, of which the switching will only be a minor part. Main problem will be interfacing hosts to X25 network. It is felt that MIDNET will not require a regional switch but as an X25 NET, will be ready necessary.

SWUCN

Presenter: J. Thomas

Current configuration

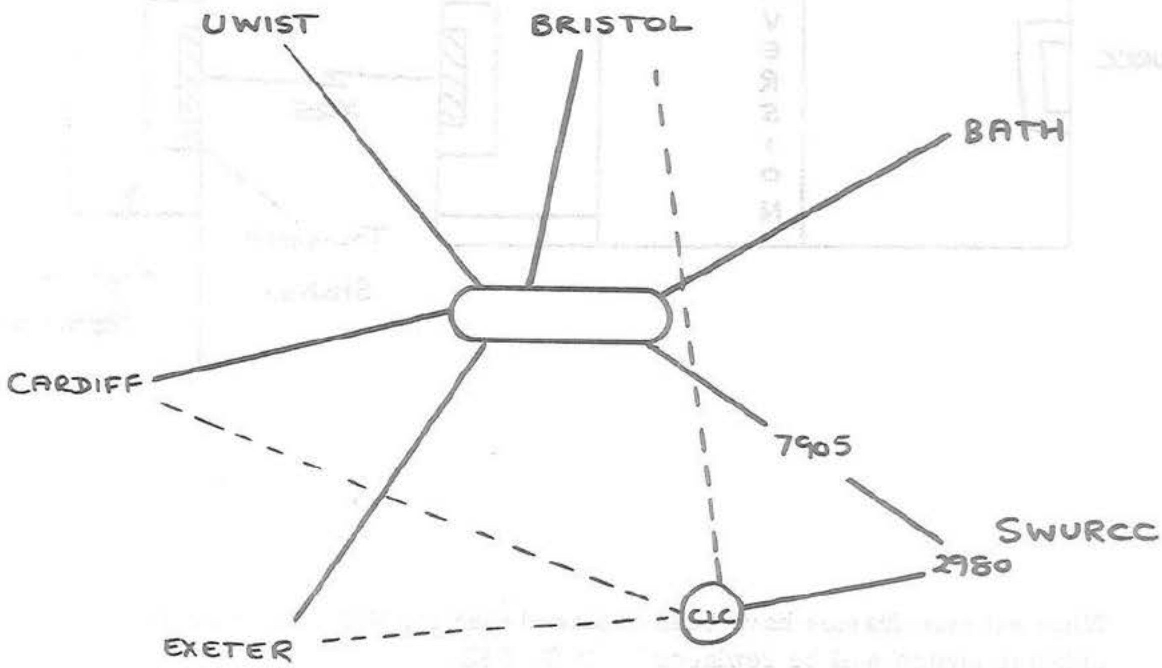

Dotted links are to provide access to 2980 if the switch or lines should break.

Running about 10,000 terminal connections/month and   jobs/month.

Problems:

a. existing protocols - outdated
   - SWUCN dependent

b. need to replace system 4 mainframes

c. aging central switch - network vulnerable to failure of this machine.

New network being planned with following objectives:

a. use modern, standard protocols,

b. easier connection of new mainframes,

c. eventual replacement of central switch (by P.O. supplied central switch).

New mainframes will interface to existing network by a network Interface processor (NIP) (GEC 4070) translating X25 to SWUCN protocol.

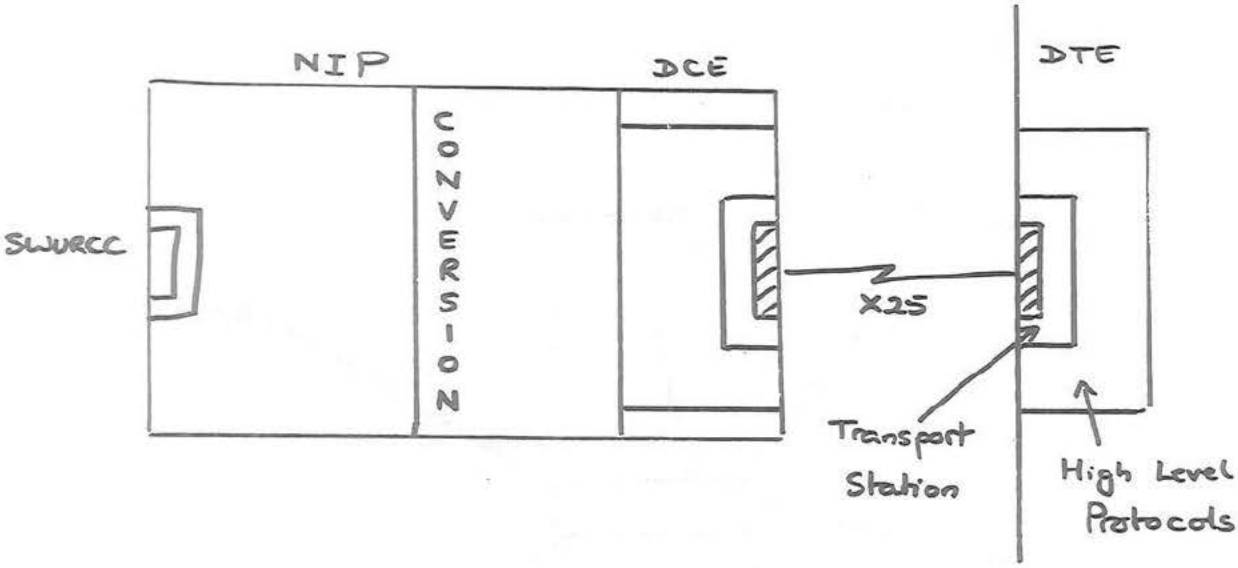

When all mainframes have been replaced then the NIP convertor and existing switch will be replaced by P.O. PSE.

Development plan

a. Instal the 4070 NIP and link to SWURCC.

b. Write and test FTP via EPSS.

c. Develop transport station and ITP.

d. Write FTP, TS and ITP on new mainframe.

e. Mount X25 and instal production NIPs.

f. Integrate with Network.

g. Further protocol development and job management.

Expect X25 software from GEC at end of year. The NIP has already passed messages to existing network.

Protocols in use will be:

FTP - Study group 2
Transport Station - simplified INWG
Terminal - X29 between NIP and mainframe.

Liverpool NETWORKSHOP 2

The SOUTH WEST UNIVERSITIES COMPUTER NETWORK

Our present protocols were developed locally many years before the CCITT started thinking about X25. They depend to greater or lesser extent on the local hardware - System 4's and CDC 1700, and System 4 operating system features.

Maintenance and support of locally developed S/W can be a problem in the longer term eg staff leave.

We have to bear the full cost of attaching new computers to the Network by reimplementing our protocols (unless of course we want to add more System 4's). We're currently in the process of replacing the Bath and Bristol mainframes.

The CDC 1700 central switching computer is nearly 10 years old - and must collapse at some point - how can we replace it.

Our objectives were simple:-

(i) To use modern standard protocols, X25 and HLP's

(ii) Easy connection of new mainframes

(iii) Eventual replacement of CDC 1700 - by perhaps a regional PSE?

Our strategy is also simple:-

As each System 4 is replaced, the new mainframe will be expected to talk to the network using X25 (supplied by the manufacturer) plus transport station and high level protocols. FIP for files and X29 for terminals - in the absence of an agreed VTP.

The new machines will be front ended by a NIP which will convert the protocols implemented on the mainframe to the present SWUCN protocols.

From the 1700, the NIP will look like a SWUCN System 4 computer. The mainframe will see a DCE when it looks towards the NIP.

In the long term, when every mainframe has a NIP, we could replace the CDC 1700 by a regional PSE, ditch all the protocol conversion software in the NIP and use them to link in campus networks, terminal concentrators and RJE stations.

To bring all this to fruition we have set up a NIP Project

Phase one started last September when we installed a development configuration NIP - a GEC 4070. Since then we have converted much of our CORAL software - written for the 7905 - and were able to talk to the CDC 1700 last December. We have since talked through the 1700 to other sites and phase one is virtually complete.

Phase two entails the implementation of FTP on the NIP and the necessary software to convert from SWUCN FTP to the new FTP. To test this out we are connecting to EPSS - using Harwell written, GEC distributed software.

In Phase three, which is just beginning, we will write a Transport Station and ITP for the NIP. The TS is simplified INWG 96, the ITP will be X29.

Phase four carried out in parallel with two and three concerned with mainframe software.- TS, FTP, ITP, the manufacturer will be supplying X25.

Later this year, having checked out FTP via EPSS, we will remove the EPSS software and replace it with X25 to be provided by GEC. We will also purchase and install cut down production configuration NIP's for Bristol and Bath to go with their mainframes.

Phase six covers the actual link up between NIP and mainframe which we optimistically aim to be complete in time to provide an experimental user service by October 1979.

Beyond that date further protocols, eg MAIL and decent Job Management software will be produced.

I will be surprised if we manage it all for less than 10 man years effort.

There will be about 12 months overlap period when old and new machines run side by side. We will then use TDM's to split the single 48k links to the CDC 1700 into two channels.

SOUTH WEST UNIVERSITIES' COMPUTER NETWORK 1977

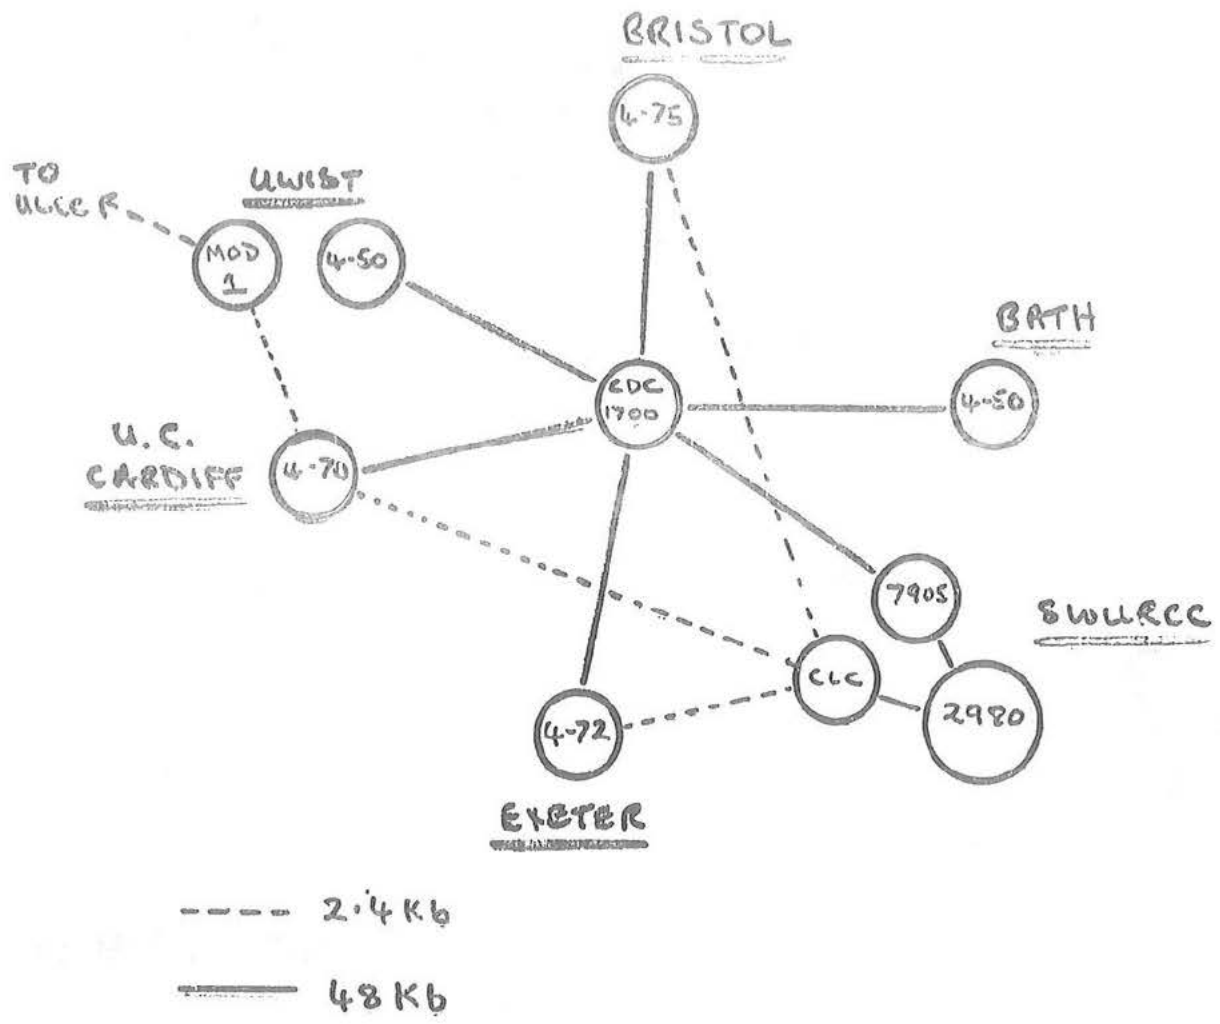

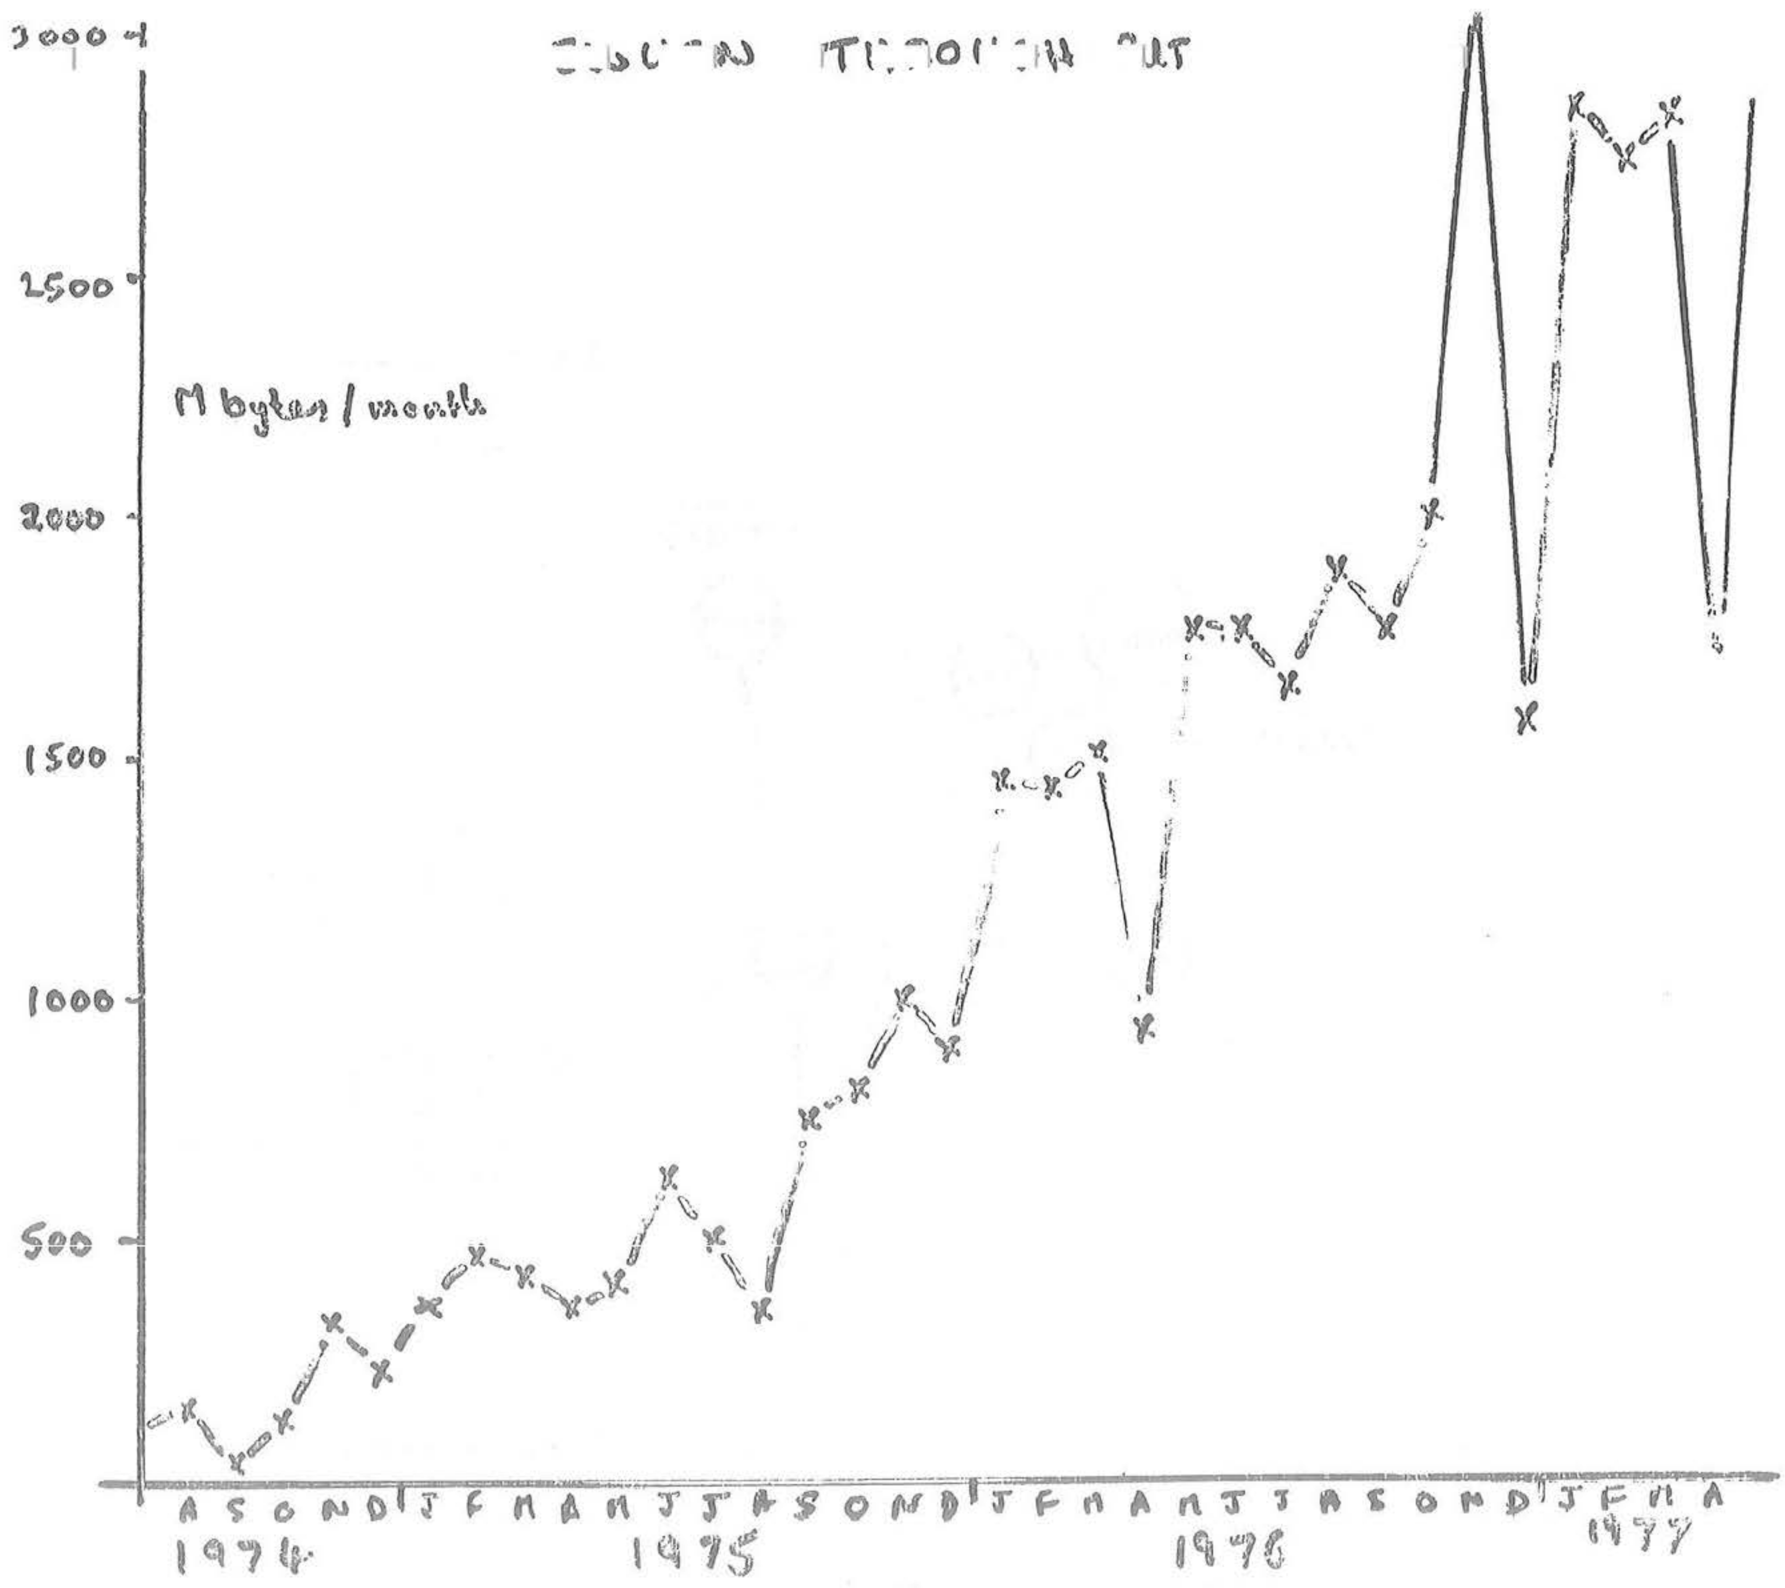
3000
2500
2000
1500
1000
500
M bytes / month
A S O N D J F M A M J J A S O N D J F M A M J J A S O N D J F M A
1974
1975
1976
1977

SWUCN DEVELOPMENT

PROBLEMS

* EXISTING PROTOCOLS - OUTDATED
  - SWUCN DEPENDANT

* NEED TO REPLACE SYSTEM 4 MAINFRAMES

* AGEING CENTRAL SWITCH (CDC 1700)

OBJECTIVES

* USE MODERN, STANDARD PROTOCOLS

* EASIER CONNECTION OF NEW MAINFRAMES

* EVENTUAL REPLACEMENT OF CENTRAL SWITCH

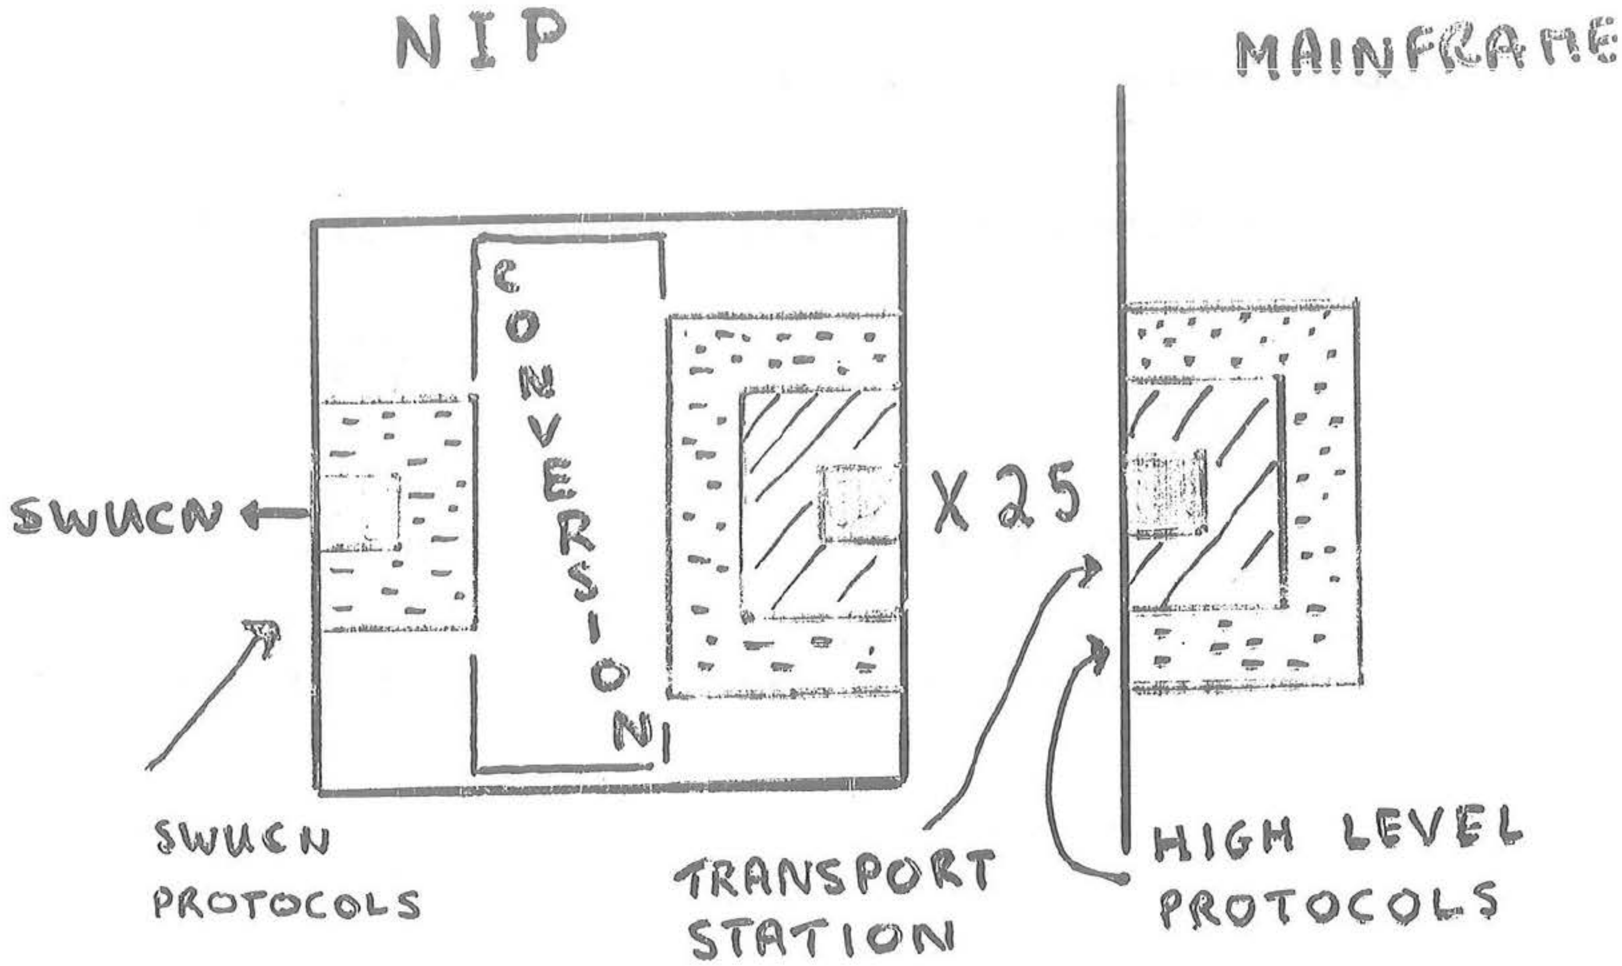
NIP
MAINFRAME
CONVERSION
X 25
SWUCN
SWUCN PROTOCOLS
TRANSPORT STATION
HIGH LEVEL PROTOCOLS

SWUCN DEVELOPMENT

NIP PROJECT - OVERVIEW

INSTALL DEVPT. NIP + LINK TO SWUCN

WRITE + TEST FTP VIA EPSS

DEVELOP TRANSPORT STN. + ITP

WRITE FTP, TS + ITP ON NEW MAINFRAME

MOUNT X25 + INSTALL PRODUCTION NIPs

INTEGRATE WITH NETWORK - EUS

FURTHER PROTOCOL DEVPT + JOB MANAGEMENT

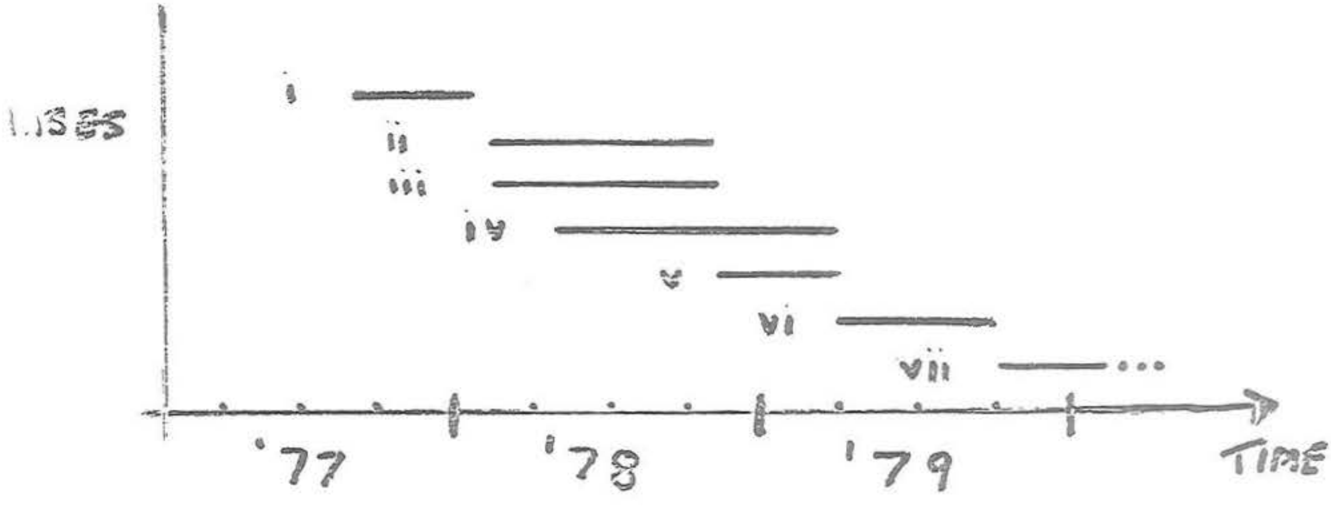

S.R.C.

Presenter: A. Peatfield

Current configuration

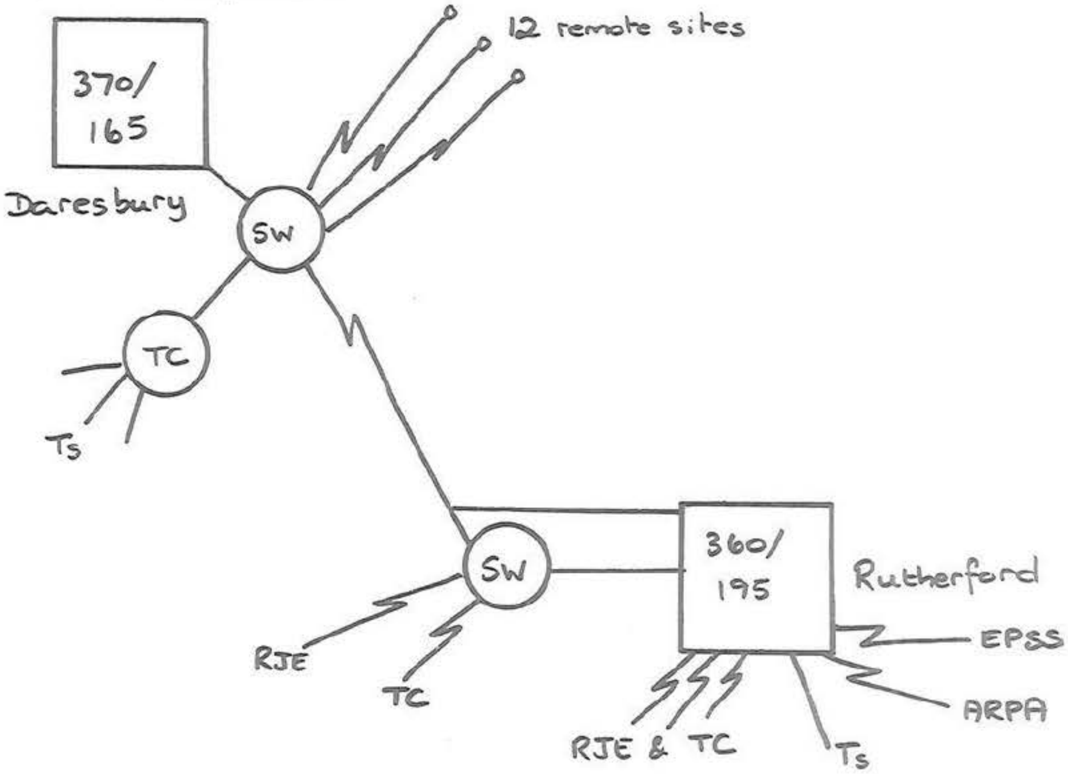

The network runs EPSS like protocols. About 300 terminals connected and 70 work stations.

Going to move to X25 and FTP as quickly as possible.

SRC ICF Network

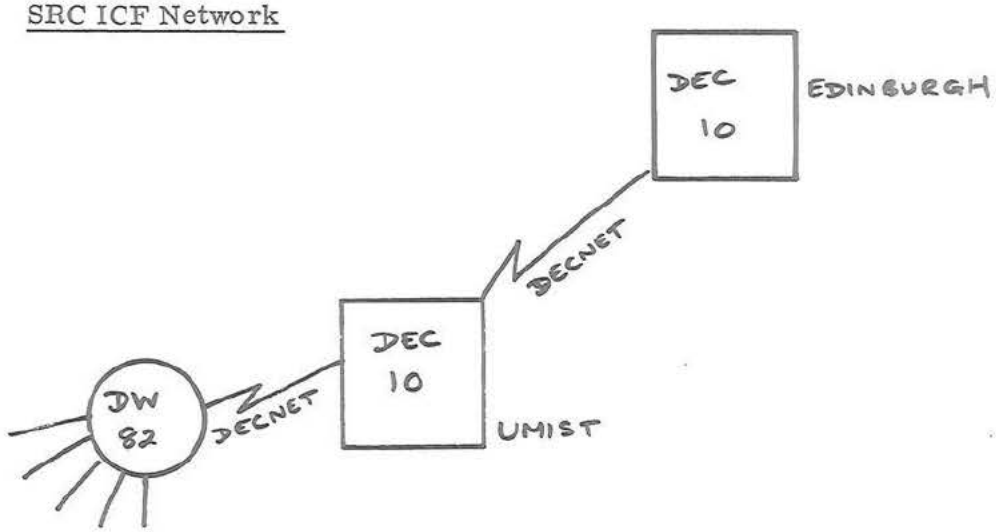

At present no interconnection between two networks. Study being produced (by Hatfield Poly and Essex University) as to how to adapt DECNET for X25.

Status and logging facilities recently added.

Report on the North West Network Service

John Lindley

NW Network Service Coordinator

1. Introduction

The North West Universities network is now operational and allows the shared use of the mainframes at Lancaster, Liverpool and Salford via keyboard terminals and batch jobs to users at those three universities and also to users at Keele University. The aim of the paper is to report briefly on the service, on what it is, on how it compares with what was planned and on some of the problems experienced so far.

2. History

Regional cooperation between the North West computer centres has its origins in the Flowers' Report and the consequent establishment of the Regional Centre at Manchester, of which the four other North West Universities of Keele, Lancaster, Liverpool and Salford were the first external users. Access to the UMRCC 1906A/7600 systems is obtained via 7020 batch RJE terminals or via CTL Satellites emulating 7020's. In 1974 the Computer Board awarded an ICL 1906S machine to Liverpool on condition that something like 25% of its power was made available to Keele, Lancaster and Salford. Again, to achieve use of this power, a star network of 7020 style connections centred on Liverpool was set up (Figure 1).

The shortcomings of these multiplying star RJE networks was early realised, with their requirements of dedicated equipment and more particularly, their dearth of keyboard terminal access (one dedicated and poorly served MOP terminal per site to Manchester and dial-up access only to Liverpool).

It seemed that there would be long term economic advantages in the development of a distributed network linking the five sites and offering easy access from each site to all mainframes in the region. More importantly, such a network would be a means of greatly assisting the growing collaboration between the five universities which was seen as the only way of attempting to satisfy the ever increasing demands of the region's computer users for power, special peripheral facilities, and a wide range of programming languages and application packages.

Particular aspects of this collaboration which the network was expected to assist were:

(i) Resource sharing;

(ii) Package specialisation;

(iii) Expertise sharing;

(iv) Future flexibility in machine provision;

(v) File transfer

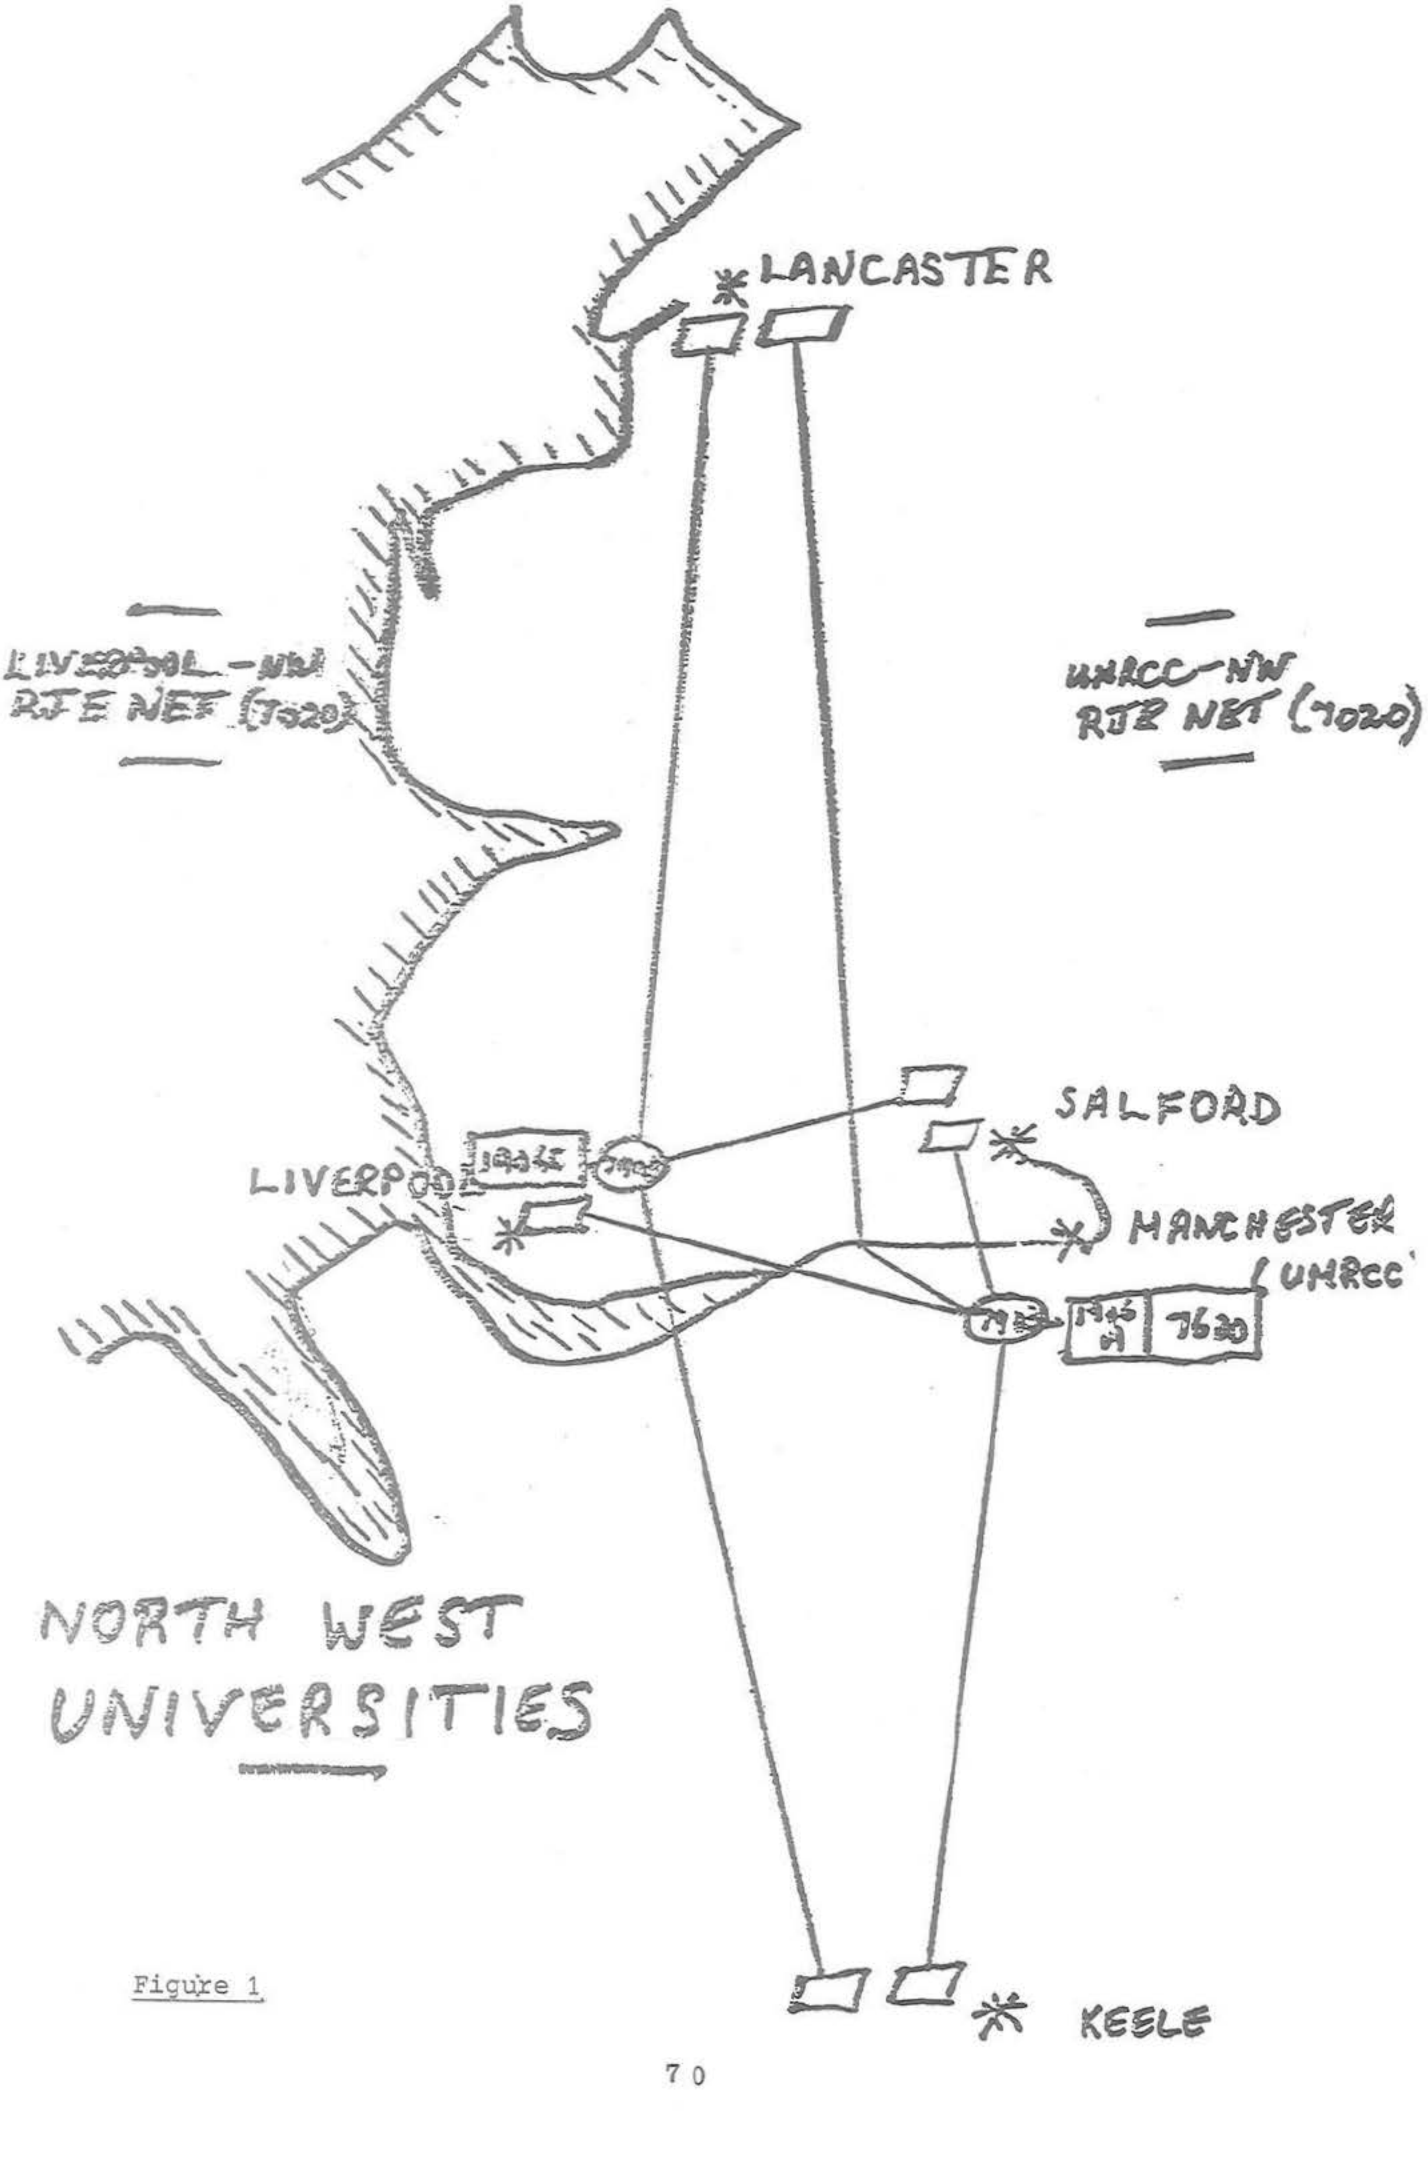
LANCASTER
LIVERPOOL-NW
RJE NET (7020)
UMRCC-NW
RJE NET (7020)
SALFORD
LIVERPOOL
MANCHESTER
(UMRCC
7600
NORTH WEST
UNIVERSITIES
KEELE

Figure 1

3. The Plan

A Network Working Party representing all five universities was set up in 1974 to examine the feasibility of setting up such a network and to make recommendations. Its conclusions were agreed by the respective university Computer Committees and successfully presented as a networking proposal to the Computer Board in May 1976.

The required network facilities were to include:

1. Remote Job Entry — to any site from any network batch RJE terminal;
2. Interactive Working — to any local or remote site from any network keyboard terminal with all the facilities of the mainframe accessed being available;
3. File Transfer — between any two mainframes initiated from any keyboard or RJE terminal;
4. Accounting and monitoring of network traffic and performance;
5. Packet switching
6. Gateway potential (initially via EPSS)

It was essential that the network should use as much of the existing hardware in the region as possible (mainly ICL 1900 mainframes with ICL 7905 front ends), but be capable of ready connection to new types of mainframes as those were introduced into the universities. It was highly desirable, in view of the limited amount of spare programming capacity and communications expertise available in the region, that the network should be developed from an existing design and existing software. Of several schemes examined, the RCO network centred on Edinburgh and the ICL/DTI sponsored GANNET project appeared to be worthy of detailed consideration, and ultimately GANNET was chosen and the software and much expertise made available by ICL. The GANNET project had commenced in 1972, and after many man years of effort was in a working form operating between two Government sites (at Norwich and Darlington), which was successfully demonstrated to the universities in late 1975 and early 1976. While the working network was based on ICL 7905's operating as Network Interface Processors (NIP's) and allowing connection of ICL 1900 series Host computers operating under George 3/4, the design was modular and intended to be essentially independent of the particular type of NIP or HOST used. At the time the NIPs only supported teletype terminal connections and job and file transfer and file listing facilities between hosts.

A phased development of the network (Figure 2) was envisaged as follow:

Phase I — (Target date late 1976, actually achieved July 1977) Three sites network linking Lancaster, Liverpool and Salford and offering similar facilities to those of the Norwich/Darlington connection, with HDLC style communication between NIPs at 4800/9600 baud.

PROPOSED NETWORK

3 SITES INITIALLY.

5 SITES

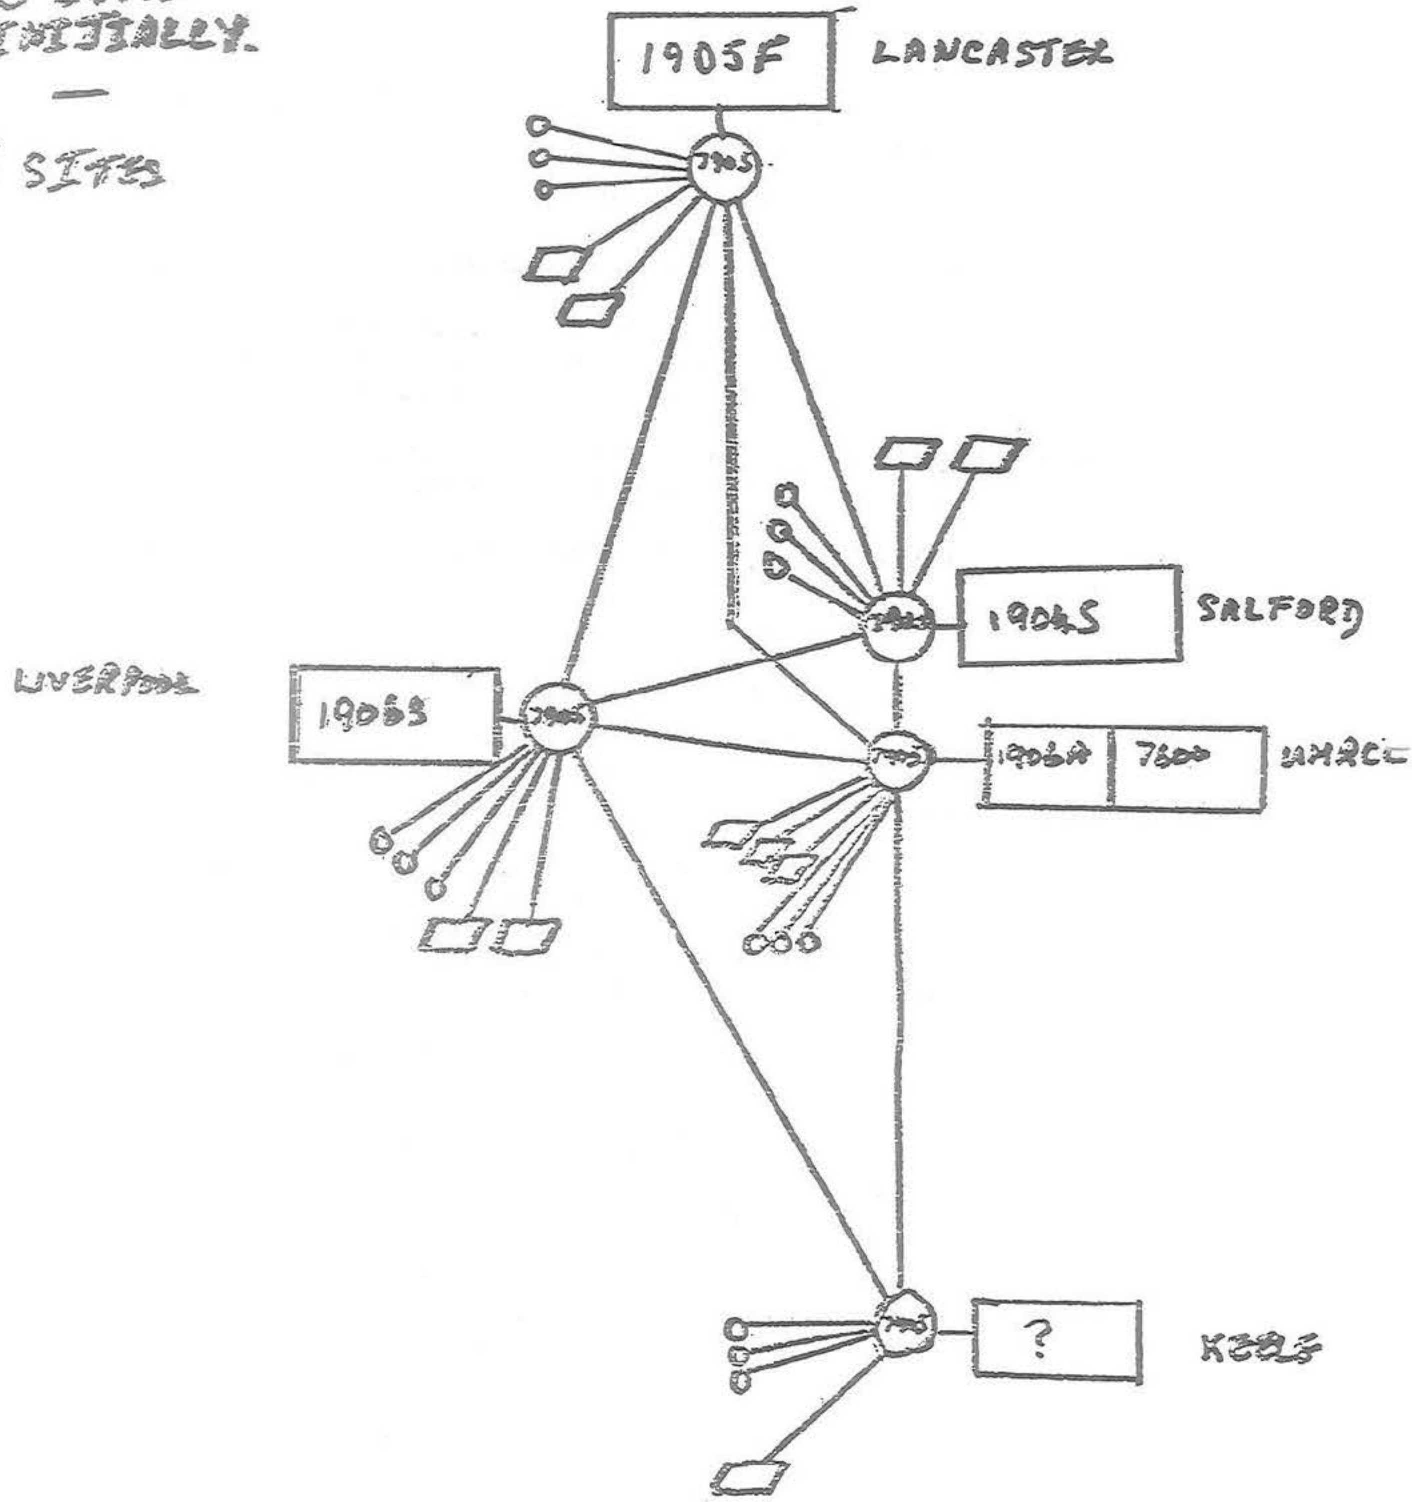

Figure 2

Phase 2 (Target date September 1977)
Support by NIPs of RJE terminals and keyboard terminal concentrators;

Relaxation of Username restrictions to allow controlled use of the network by the large number of users normal in a University community.

Phase 3 (Target date September 1978)
Connection of Manchester to network;
Support by NIP's of mini-hosts (intelligent RJE's etc).

Phase 4 Connection of a new host at Keele to network;
EPSS gateway.

4. Present Network and Service

As indicated above, Phase 1 was completed essentially by July 1978, although the resilience of the hardware and software and certain aspects of the user interface were considered inadequate to offer a service immediately. Accordingly it was only in October last that an initial and rather tentative service began to be offered on an evening only basis. Resilience and the ability of the network to handle sufficient keyboard terminals are vital to Lancaster and Salford since the 7905 is essentially their only terminal handling front-end.

By January of this year the purchase of a 7905 NIP for Keele had been approved and the machine installed, thus offering access for Keele users to the network (Figure 3). At the same time the hours of service offered at Liverpool and Keele were extended to cover the latter half of the prime shift period. Currently the hours of service offered at Lancaster are being similarly extended, and Salford is likely to follow suit in the immediate future. The aim is to operate the network at all sites from about 10.15 a.m daily, leaving a common early morning period for hardware and software maintenance, and software development.

Aspects of the current network and service include:

(i) Current line speeds are all 4.8Kb only;

(ii) In most cases lines are shared with the existing Liverpool RJE star network (until Gannet is operational full-time) using Racal Milgo 5500/96 modems with dual porting capability (Figure 4).

(iii) No attempt has yet been made to connect 7020 RJE terminals to the NIPs, mainly since this would.imply a need to undertake upgrading of 7905's to 7906's, and because the three existing hosts have alternative front-ends to which their RJE terminals are connected and through which users may still carry out network operations. A review of this requirement is currently being urgently undertaken;

(iv) Figure 5 lists briefly the main user commands available with the network. Essentially the keyboard terminal user first logs on to the network (using a numerical username). He then has a number of NIP-based facilities available to him including the ability to interrogate a global NOTICEBOARD and his personal MAILBOX (which will include information on any files and jobs that have been transferred over the network for him). To use any

ACTUAL NETWORK

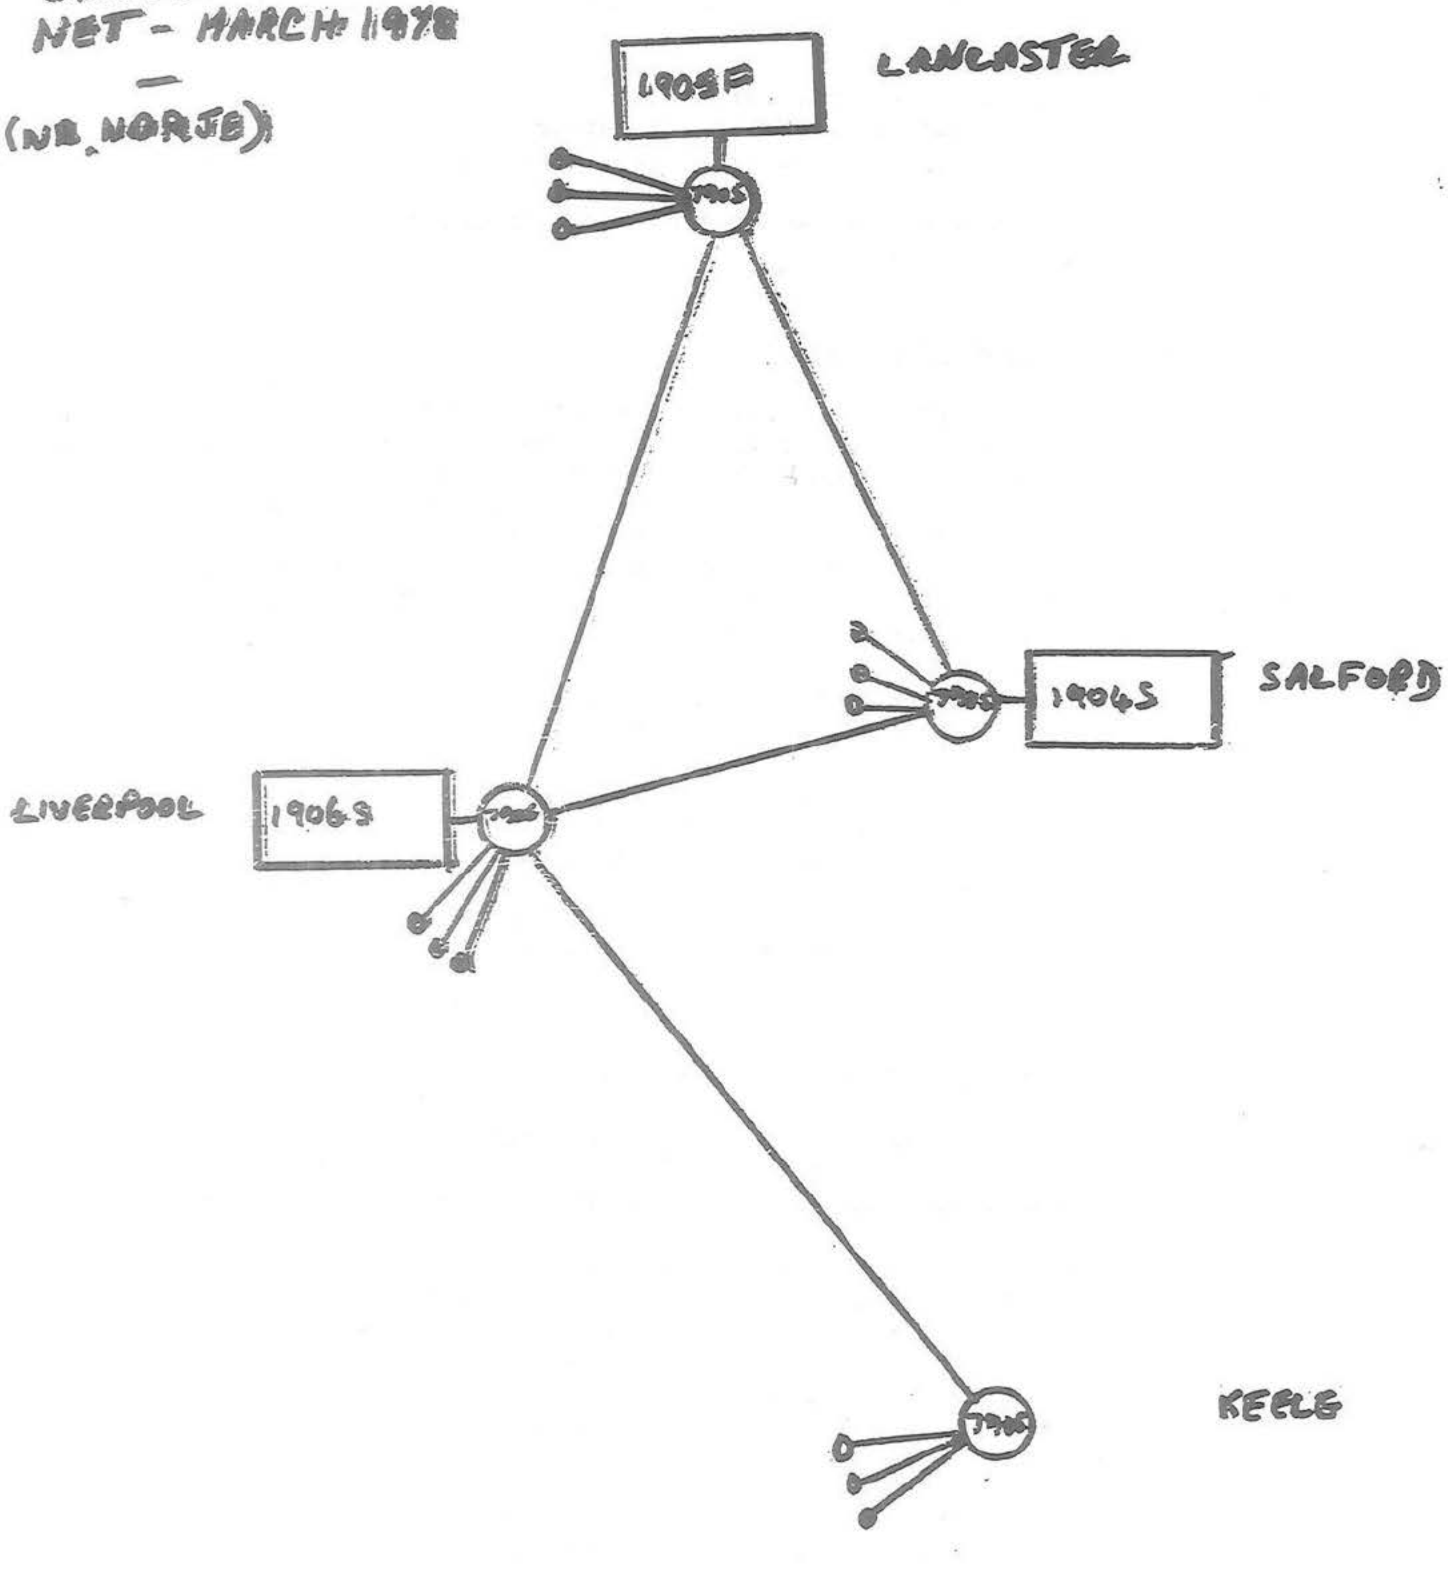

Figure 3

ACTUAL NETWORK

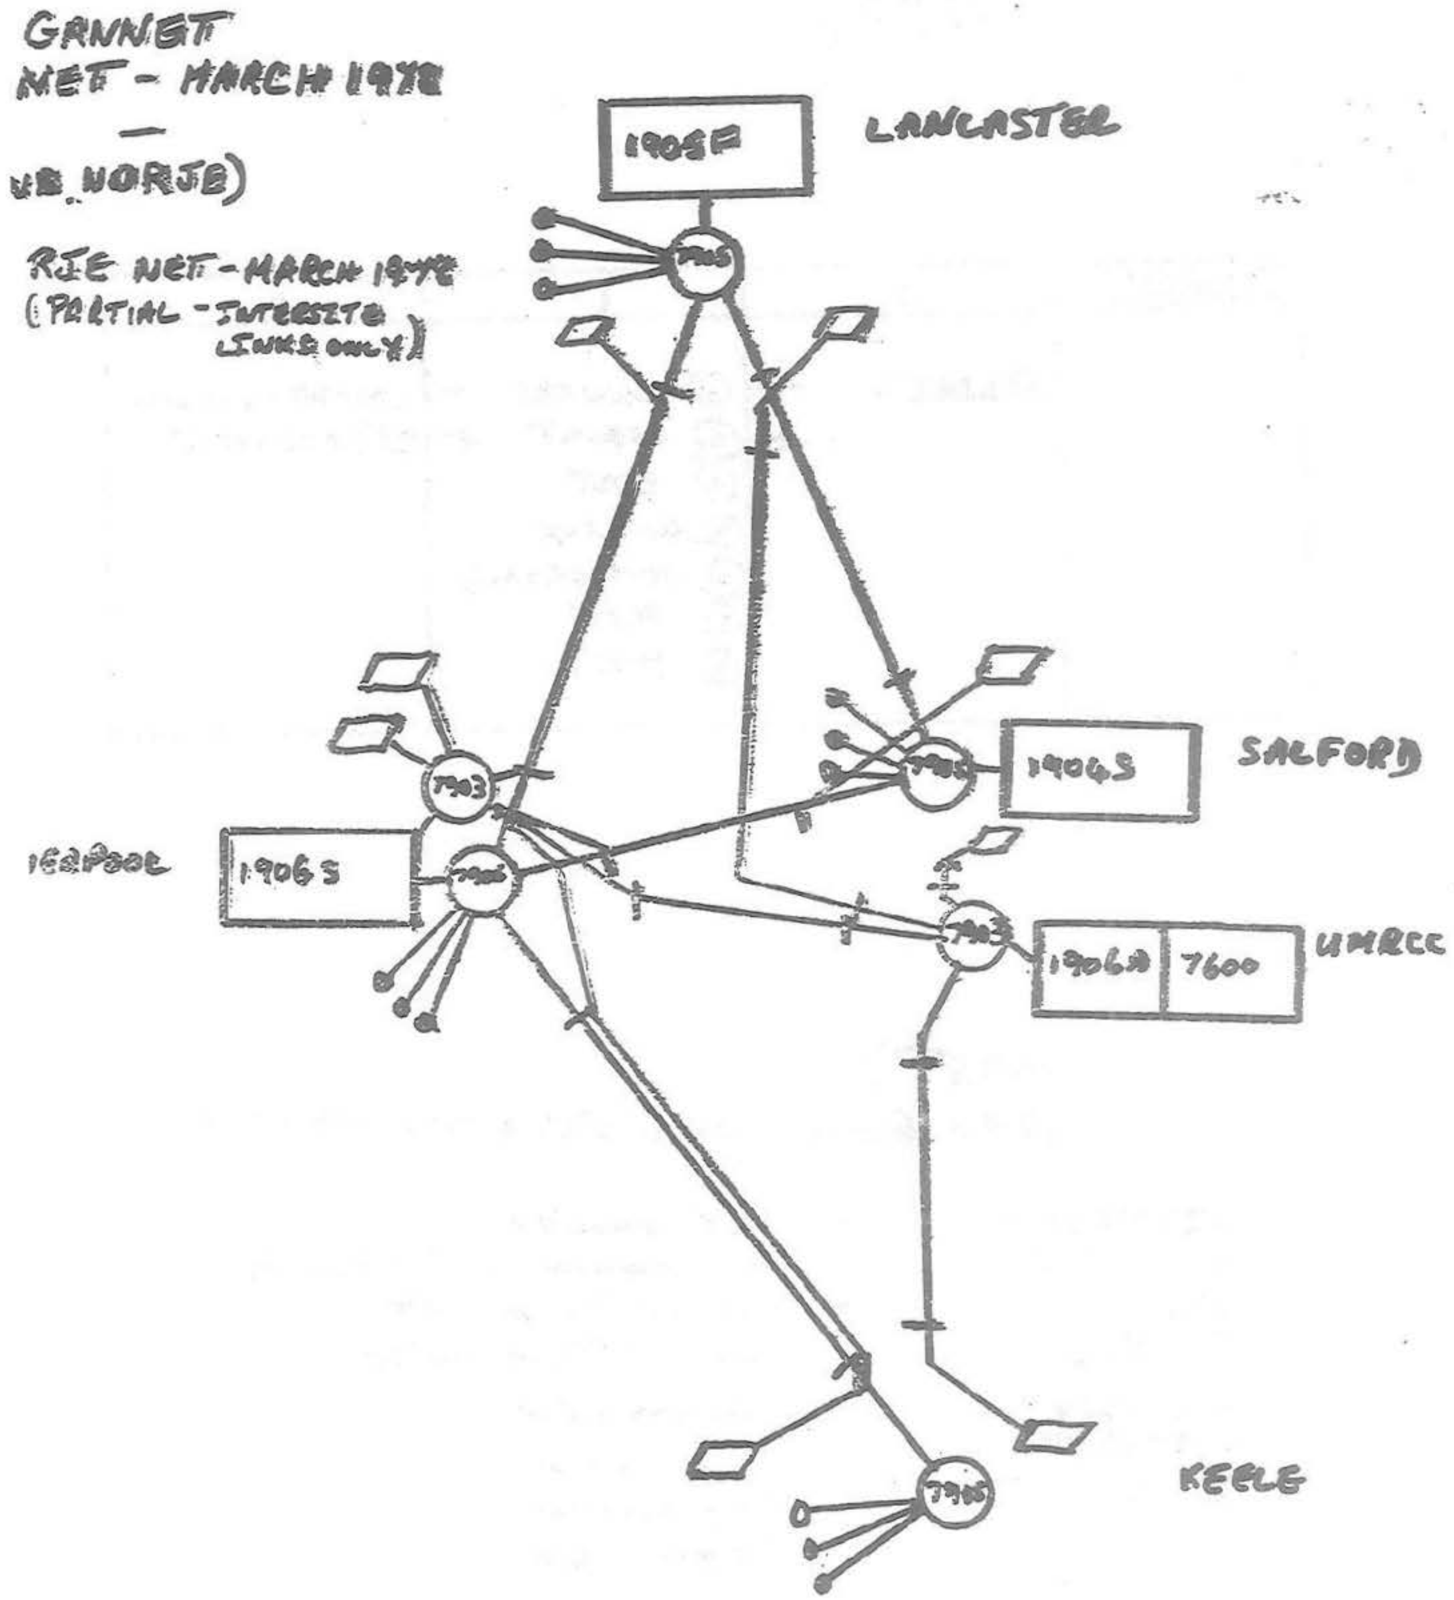

Figure 4

NETWORK USER COMMANDS
(NIP)

Level / Commands	0	1	2
	@LOGIN	@ CONNECT @ LOGOUT @ EDIT @ MAILBOX @ NOTICEBOARD @ POST @ STATUS	GEORGE commands @DISCONNECT

(HOST)
1900 George commands with network features.

LISTFILE	–	HOST parameter
TRANSFER	–	New command – File transfer
JOB	–	HOST and PR parameters
RUNJOB	–	HOST and PR parameters
ABANDON	–	HOST parameter
ATTRIBUTE	–	PR parameter
DISCONNECT	–	PR parameter
LOGIN	–	PR parameter
WHATSTATE	–	
NEWSTAT	–	New command – Status of command well
WHATNWCOM	–	New command – Number of commands in well
NEWLIST	–	New command – List of commands in well
NEWABANDON	–	New command – Abandon welled command
MSGNETWORK	–	New command – Status of hosts in network

Figure 5

Host on the network he must first CONNECT to that Host, quoting the name of the Host, and he then operates exactly as a conventional local user. With 1900 Hosts the main network facilities include the TRANSFER command for file transfer, and extensions to the George, JOB, RUNJOB and LISTFILE commands to enable transfer of jobs and files for output to other Hosts. Figure 6 shows a typical brief user session with explanation.

(iv) The basic hooks exist in the network software for monitoring availability and reliability and for accounting for usage of the network. Basic reporting software to analyse this information is still being written so that no automatically recorded information is yet available in a meaningful format. In the meantime manual recording of availability and reliability is being undertaken (which shows a steady improvement in recent weeks towards the target of not more than one NIP break per week per site - some 100% weeks have already been recorded), while existing site host accounts programs give some guide on network usage. Figure 7 gives some figures for usage of the Liverpool 1906S from the region during February and March of this year. In addition to the recorded figures very substantial use (1500 jobs per month) is made by Keele of the Liverpool cafeteria service.

(v) Present usage of the network, in addition to those sharing Liverpool 1906S for its power, include use of specialist packages such as Genesys, use of particular compilers such as Algol 68 and Basic and of special facilities like the microfilm service provided through Liverpool.

(vi) Users and operators now have available to them full manuals on the network facilities and operational procedures.

(vii) Work on connecting UMRCC into the network is under way with a new target date for a service of January 1979. The new Host at Keele, a GEC 4082, will shortly be delivered and work will begin on linking this to the network, again hopefully to be completed by early 1979 (Figure 8).

Note

```
break in
< >@LOGIN:7                                                    (a)
LOGGED IN ON WEDNESDAY 02/02/77 AT 15.13.21
< >@STATUS:LIVERPOOL                                           (b)
LIVERPOOL ONLINE AND AVAILABLE FOR MOP
< >@CONNECT:LIVERPOOL                                          (c)
CONNECTON OK
LIVERPOOL UNIVERSITY GANNET G4 MARK 8.52/3                     (d)
15.15.10←LN TESTJOB,:NW03,PR LANCASTER
TYPE PASSWORD←!!*
STARTED :NW03,TESTJOB, 2FEB77·15.15.25 TYPE:MOP
15.15.46←ED PROGFILE                                           (e)

EDITOR IS READY

0.0← TS/DIMENSION/,.E,I/,A(15)/
8.29← #53,R/20/10/
53.17← E
15.17.10←LF PROGFILE,FRS/DIMENSION/,LI1                        (f)
      DIMENSION IA(20),IB(20),A(15)
15.17.21←LF PROGFILE,*LP,NU                                    (g)
15.17.25←LF PROGFILE,*LP,HOST SALFORD                          (h)
15.17.27←AU PR SALFORD                                         (i)
15.17.33←TF TO PROGFILE,HOST SALFORD                           (j)
15.18.05←RJ TESTRUN,TESTJDF,JD(JT 50,MZ 40K)                   (k)
15.18.19←LT                                                    (l)
MAXIMUM ONLINE BS USED 13KWORDS
15.18.26←0.02 FINISHED:1 LISTFILES HERE,1 LISTFILES ELSEWHERE(m)
< >@DISCONNECT                                                 (n)
DISCONNECTION OK
< >@MAILBOX                                                    (o)
01        02/02/77        15.17
PROGFILE TRANSFERRED FROM LIVERPOOL TO SALFORD
END OF MAIL - DELETE ? A
< >@STATUS:SALFORD                                             (p)
SALFORD ONLINE SUPPORTING MOP BUT NO CONNECTIONS
< >@LOGOUT                                                     (q)
```

Figure 6 (i)

Notes:

(a) The user *logs in* to the *network* with a *network username* of 7.

(b) The user ascertains whether he can run a MOP session at Liverpool.

(c) Liverpool is available for MOP so the user @CONNECTs to Liverpool.

(d) The user *logs in* to GEORGE 4 at Liverpool using a PR LANCASTER parameter so that all output will be routed to Lancaster automatically, unless explicitly routed elsewhere.

(e) The file PROGFILE held at Liverpool is EDITed.

(f) One of the amended records in PROGFILE is listed to the terminal.

(g) The file PROGFILE is listed to a lineprinter at Lancaster.

(h) The file PROGFILE is listed to a lineprinter at Salford.

(i) The user changes the default destination of output resulting from LF or RJ commands from Lancaster to Salford.

(j) The file PROGFILE at Liverpool is copied to a file PROGFILE at Salford.

(k) A job TESTRUN using the file TESTJDF is submitted to Liverpool. Output from this job will be to Salford unless it is explicitly routed elsewhere from within the job.

(l) The user *logs out* of GEORGE 4 at Liverpool.

(m) The 1 LISTFILE HERE refers to (g) and the 1 LISTFILES ELSEWHERE refers to (h).

Figure 6 (ii)

NETWORK USAGE

(JOBS SUBMITTED TO LIVERPOOL)

FEBRUARY 1978

	Keele	Lancaster	Salford
Terminal Sessions	101	107	40
RJE/Remote Jobs	227	17	52
Total	328	124	92

MARCH 1978

	Keele	Lancaster	Salford
Terminal Sessions	227	258	27
RJE/Remote Jobs	430	95	62
Total	657	353	89

Figure 7

ACTUAL NETWORK

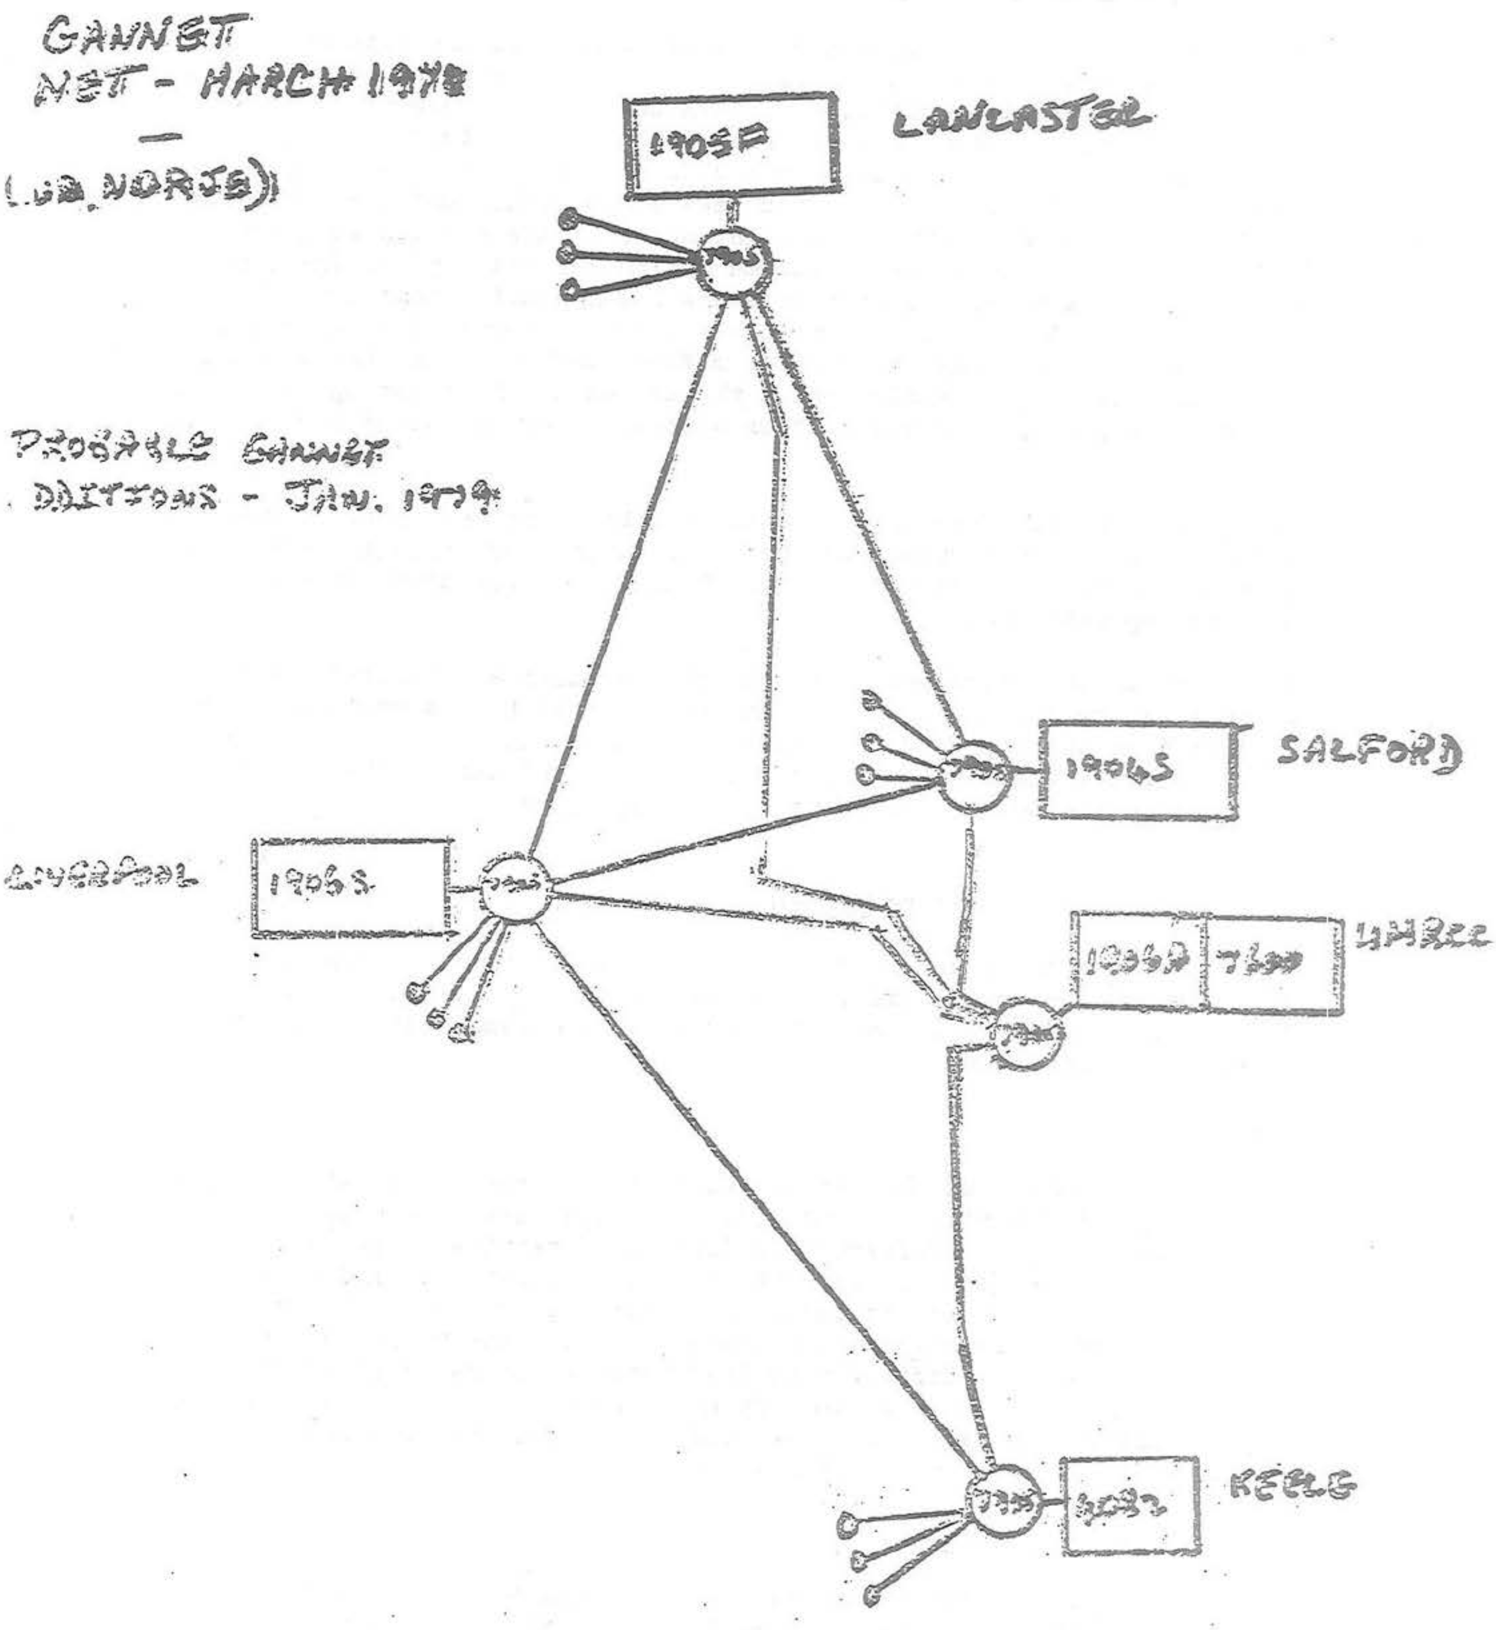

Figure 8

5. Management Aspects

The aim, as regards the direction and management of the network, and of North West collaboration in general, has been to keep the structures as simple and informal as possible. Thus there is no committee above the Directors, each of whom is responsible to his own university policy committee. Figure 9 indicates the current committee structure, showing the existence of the Joint Project Team responsible, under the Project Leader, for the development and maintenance of the network software. Collaboration in the areas of documentation has taken place for some time, hence the existence of a separate Documentation Committee. The Operations and Applications Committees are more recent developments, the former coordinating the service offered and involving the manufacturer's engineers and support staff, which the latter has been set up to examine in detail opportunities for further software specialisation and expertise sharing.

The Network Coordinator (Dr J D Rice) has been the key figure responsible overall for the development of the network, and more recently a Service Coordinator role was conceived to look after the operation of the network and its exploitation.

Whenever possible problems with network operation are handled locally; responsibility for the diagnosis of the cause of breaks rests with the Project Team and any truculent hardware reliability problems are handled by the Service Coordinator and and ICL CED Coordinator. Daily checks are made by the Service Coordinator with all sites.

6. Some problems experienced

Many problems have been experienced in the development of the network and in establishing the present initial service. Although some have been technical, probably the majority have been human ones. Major aspects have included:

(i) Project

Development of the network software has been seriously hindered by staff mobility (communications expertise is a highly saleable commodity). Differing conditions of service of staff at different sites have caused problems (e.g. over travelling and subsistence expenses). It has often been difficult to make all machines in the region available simultaneously for development sessions and to find sufficient trained manpower to man them at these times. Clashes of loyalty have occurred since it has not been possible totally to divest most members of the project team of support responsibilities at their home site.

(ii) Decision making

The participating universities have wished to maintain their independence as far as possible. It has not always been clear who finally decides on aspects of policy and on more detailed technical points with regard to the development and exploitation of the network. Thus there has been a degree of inertia and sometimes of backtracking over making certain decisions. The role of UMRCC as a national centre with national responsibilities has meant that Manchester's involvement has had to be a particularly cautious one.

COMMITTEE STRUCTURE

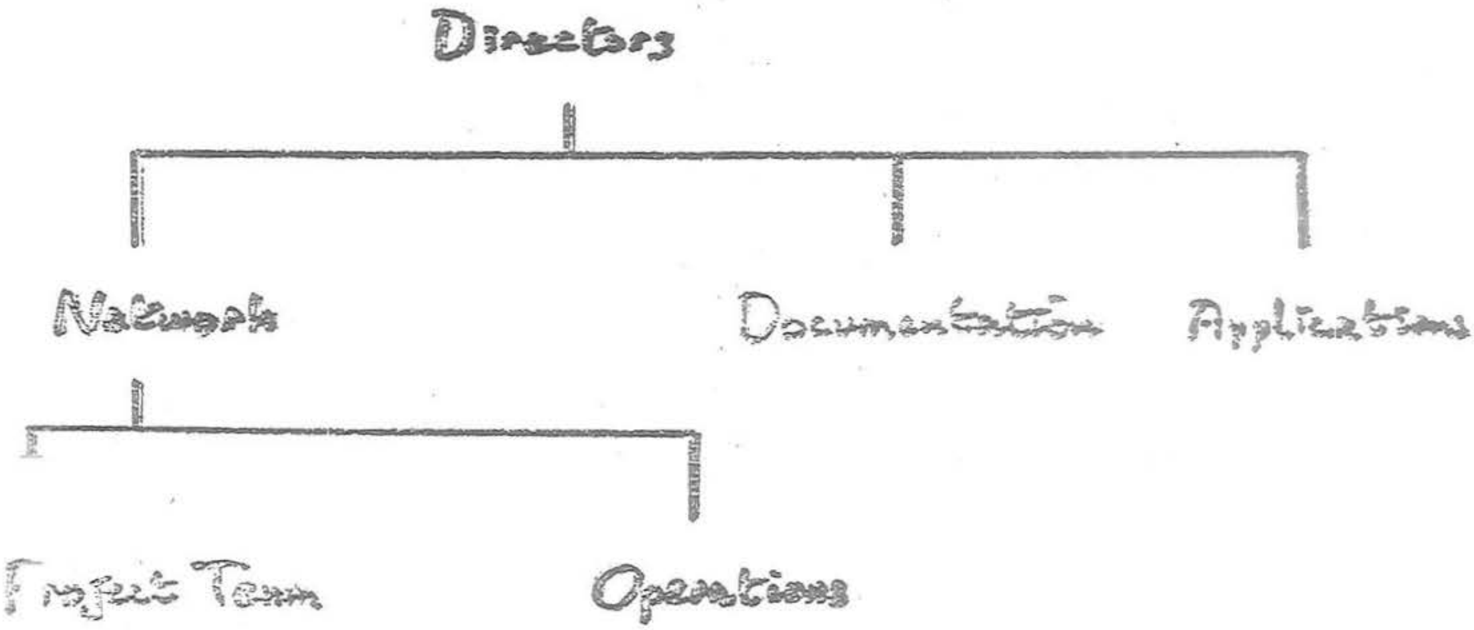

\+ Network Coordinator
Network Service Coordinator
Project Leader

Figure 9

(iii) Technical

Some technical problems have existed for some time and are as yet not finally resolved. Problems with regard to Usernames and the connection to the NIP's of RJE terminals have been mentioned above and solutions are being actively pursued. Some disagreement over aspects of the user interface to the network are still being resolved. There is a continuing danger, with developing protocol standards, of the network project having a continually moving target. Currently there seems to be advantages in using FTP for file transfer, which are under investigation, but the question of whether to use the full X25 protocol has been shelved for the time being.

(iv) Operational

The main problems are in maintaining adequate human coordination between sites, and in getting full operator understanding of the network. To this end a series of operator training sessions, involving staff from all sites, have been started.

(v) Users

There are serious problems, particularly in the present interim period, of keeping users in touch with whether and when the network is operational.

(vi) Hardware and Software Reliability

The data lines and modems have presented few reliability problems. Crucial therefore to the reliable operation of the network has been the performance of the 7905 NIP's. On the software side the position has been steadily improving although some faults require a significant load before manifesting themselves. On the hardware side there have been some serious and prolonged problems at individual sites and between sites (e.g. incompatible disc drives). Increasingly close collaboration with ICL CED will, it is hoped, gradually improve hardware reliability of the 7905's to figures constantly approaching 100%, with very rapid rectification of faults. Host reliability is regarded as a local site problem and, of course, varies from host to host.

7. Conclusion

The North West Network is operational and providing a useful service. Further development is taking place and major milestones will include whole day service and, most importantly, the connection of UMRCC and the new host at Keele. Further into the future there will be a new host (possibly ICL 2950/60) at Lancaster to connect and also possibly interactive systems at Salford and Liverpool. However development within the present project is scheduled to be concluded by March 1980. Post Office plans for a national packet switched network are eagerly awaited.

It is believed that existing problems can and will be overcome, that the network will meet the needs for which it was set up, and will evolve to take advantage of the wider networking opportunities which will become available.

CAMPUS NETWORKS

CONSIDERATIONS FOR CAMPUS NETWORKS

Some Price Comparisons

Logic — A micro-processor costs £25; with memory and simple support chips the cost of a system to drive a VDU, with no local file store, might be £200. Random access memory costs 1p per byte.

Coding — A programmer earning £5,000 p.a. and writing eight correct lines of code per day, produces code costing £2.50 per line. The line of code might occupy 100 bytes, costing £1.

Wire — Wire, clipped to a wall, costs 50p per metre, excluding the cost of the wire.

Some Consequences

Logic containing a copy of existing code, is virtually free. If, by installing some logic, we save wire, then we can be very quickly in pocket.

Logic containing code unique to that logic is prohibitively expensive, and if at all possible we should NEVER instal only one copy of any code.

Some more Price Comparisons

For many (not all) purposes a mini may be every bit as effective as a larger system and much cheaper. Obvious examples are file editing and running other simple interactive applications. A mini will be much more likely to be successful for process control or data capture. For some (not all) purposes a large mainframe may be more effective than a mini, and so be cost effective. Obvious examples are large batch jobs and (possibly) some cafeteria services. A less obvious example is the long term storage of files, where the labour intensive manual effort needed can support many more users on a large mainframe than on a mini.

All computing is labour intensive if we count (as we must) the cost of the human user; interactive computing is especially labour-intensive, and may not necessarily be the most effective (as distinct from the quickest) way of handling a problem.

And Some More Consequences

Computer systems are increasingly being tailored to suit individual applications. In future we may see a roughly three-level system

a) 'departmental' systems at a number of sites on each campus (or firm or whatever) handling

   1) terminal concentrating
   2) RJE
   3) file maintenance
   4) some interactive work

linked to

b) a 'central' system, one per campus, handling

1) some interactive work
2) batch jobs
3) long term file storage

and in turn linked to

c) a 'national' system, one per country (England, Scotland, Wales (and Ireland?) handling

1) some interactive work
2) very large batch jobs

A Final Cost Comparison and Consequence

Ten years ago computer centre costs were 80% hardware, 20% staff; at present they are 50% hardware, 50% staff; they should move to 20% hardware, 80% staff. And most of those staff should be involved in giving advice and support to users, not in maintaining and extending the operating system.

QMC CAMPUS NETWORK SURVEY

J P Brandon
Queen Mary College, London

This talk is not about something that exists; instead I shall describe the results of a survey, conducted at QMC two and a half years ago, of the requirements for a campus network. The results still stand as nothing substantial has been done to provide the network because its estimated cost was too high - about £100K.

The requirement can be stated as a list of devices that would be connected to the network:

1 A local mainframe, via its front end processor, which is also connected to the regional or national network.

2 A direct connection to the regional or national network.

3 About six mini-mainframes some of which have several keyboard terminals attached.

4 About 128 keyboard terminals connected by concentrators; these terminals would be in groups or individually in departments.

5 About eight RJE stations - even with the disappearance of card input the printers on these devices will be needed.

6 Up to 20 data aquisition machines, some of which would have their own spooling devices.

7 About 10 mini-computers controlling experiments, which need access to greater processing power at infrequent intervals.

8 Say 10 specialised systems, such as those required for undergraduate teaching, by the library, or for production of documentation.

9 Some shared peripherals, such as plotters, microfilm, magtape drives, type-setting equipment - possibly Six in number.

10 A central shared file store into which data can be put and from which data can be retrieved; such free-standing filestores were just becoming available at the time of the survey.

The total number of ports required is about 60, at speeds from 1200 baud to 100K baud; the connection to the control filestore requires even greater baud width. The peak total throughput on the network would be about 1M baud.

The requirement still exists; implementation awaits the development of a cheap versatile network architecture.

CONNECTIONS TO THE CAMPUS NETWORK

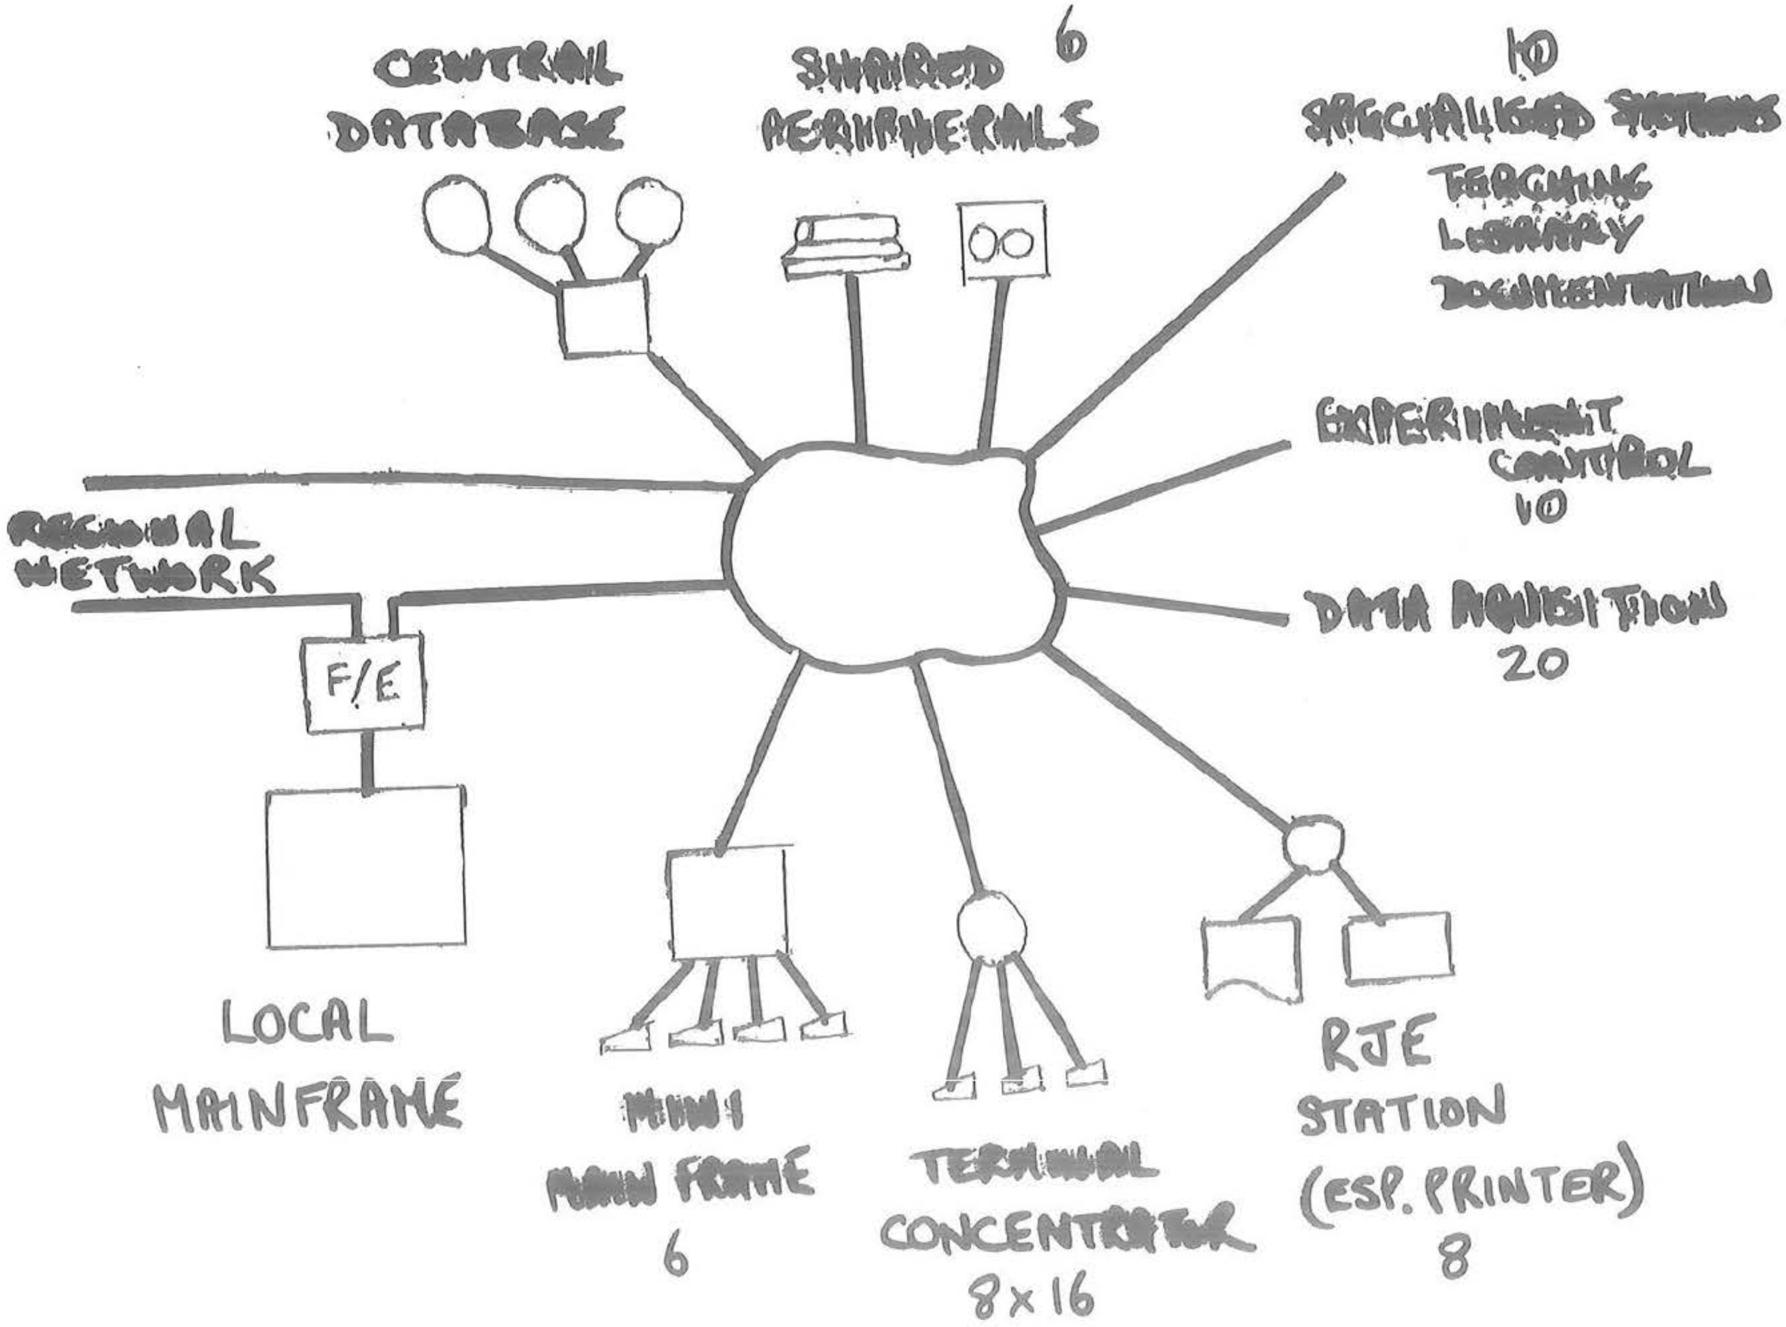

Local-Area Computer Communication Networks

Anthony West

Computer Systems Laboratory
Queen Mary College
Mile End Road
London E1 4NS

I apologise for presenting this resumé is note form, though perhaps the information is best presented thus.

Please note that my affiliation is with the Computer Systems Laboratory at QMC and not the Computer Centre.

I prefer to label the subject we are talking about **Local-Area** networks rather than Campus Networks since Campus Networks form only a small subset of networks designed for local communications. The differences lie in performance and applications suitability.

I feel that it is very important to start discussion about Local Networks for campus use from a consideration of what people want to do with them now. I also feel that the existance of a stable network within universities will encourage people to start planning communications into their problem-solutions. In other words, if you instal a network, its existence will cause the applications emphasis to shift.

Applications

I can think of many reasons (some real some esoteric) but the real-and-now reasons can be summarised as follows:-

1. Terminal Switching (multiplexing, concentrating, sharing)
2. File Transfers (data prep., RJE, data capture)
3. Job Transfer (RJE, Cafeteria, multiprocessing)
4. Sharing expensive hardware/software resources
   (e.g. COM facilities, archive storage, Graphics?)

and the more esoteric reasons are:-

5. Dynamic load balancing between machines
6. Reliable Distributed File Systems
7. Process Control of Experiments?
8. Distributed Systems Research.

Perhaps there are more reasons but I think those are the main ones. I have called those things **esoteric** for which university expertise is not yet widely available.

Network Implications

The above functions require that the network supporting them should have certain characteristics. The FTP and RJE applications require high bandwidth (several Megabaud) whereas terminal switching alone might not strain current technology too much. I use the term **bandwidth** in a loose way - basically I mean the delay generated by having to ship **B** bits of data around. The delay is the important characteristic for users of the network whereas the instantaneous line utilisation and other such factors which can also be described as bandwidth are of interest to network implementers and system managers.

Networking Spectrum

Let us now place the sort of network we are interested in in its place in the computer communications spectrum:-

	Global	**Local**	**Multicomputer**
Node Separation	> 10 Km	10 m - 10 Km	< 10 m
"Transfer Rate"	< 50 Kbaud	1 - 20 Mbaud	> 10 Mbaud
"Delay"	1 sec.	1 msec.	1 usec.
Expandable?	yes expensive (line cost)	hopefully easy & cheap (replication)	difficult

Here it is worth saying that the ease of expansion is a very important factor in these local networks. Preferably the cost of the network should increase linearly over a wide number-of-nodes spectrum.

Using microprocessors in such networks is attractive since they can be programmed to handle network control functions:-

1. Addressing
2. Flow Control
3. Congestion Control
4. Acknowledgements
5. Sequence Control
6. Packeting/Unpacketing?
7. Connection Establishment?

8. Peripheral interfacing

BEWARE - the dominant cost of using micros in a network environment will be the cost of generating the correct software. Hardware is simple to design compared to writing software for just about the most difficult of all applications - network control. The network which uses micros will only be cheap if the software in the micros serves for all micros (i.e. develop it once then simply replicate it!). The hardware is cheap **but the software is NOT.**

Characteristics of Local Networks

Local networks **should**

1. have high bandwidth
2. have a rich connection topology (perhaps even avoiding routing)
3. be cheap to interface to
4. be cheap to expand by replication of modules
5. be cheap to instal cabling for
6. have very simple network control algorithms
7. have a unit connection cost less than the unit you want to connect to the net.

Range of Possibilities

I tried to construct a list of candidates for technology which I am familiar with and which might be used for local networks. The list is (together with my subjective score out of 10 - 10 means ideally suited):-

1. X25 provided by GPO 1

I dismissed this since

a) bandwidth is low (<48 Kbaud for DATEL)
b) rather complex protocols for what is required
c) PSS is not totally under your control which
is quite a desirable attribute
d) EXPENSIVE!
e) charging is a very murky area as yet
f) bad expansion economics for more units
g) maybe you don't want virtual circuits
h) not very reconfigurable! Try moving your equipment
into another building - you have to call the
GPO engineer in to recable.

on the other hand...

a) it's off-the-shelf
b) maintained by the GPO to GPO standards
c) it will probably plug directly into your mainframes
d) having invested in X25 for local nets, you can use
the same investment for global nets.

So, there is scope for much heated argument. I prefer to avoid the issue by saying that I don't want to do distributed Microprocessor development for a micro teaching laboratory by accessing mainframe facilities via PSS!

2. Star Network based on (say) PDP11 3

This is not bad but tends to have a cost/nodecount graph that looks like a staircase. This is not necessarily a good thing depending on how your funding arrives. Problem is that the switching done by the central switch is resource consuming and the computer concerned is likely to melt under high data rate bombardment. Not a very good approach when compared to other technology (below).

3. Networks based on ETHERNET 6

This technology for local networking was pioneered by the XEROX corporation at their research centre in Palo Alto, California. There is an excellent paper about how it works:-

"Ethernet: Distributed Packet Switching for Local Computer Networks"

R.M. Metcalfe & D.R. Boggs
Communications of the A.C.M.
Vol. 19 No. 7 July 1976

That paper is well written and easy to read so I don't propose to review its content here. Suffice it to say that there is a single passive coaxial cable which connects everything tapped onto the ethernet to everything else. All stations receive everything anyone sends that they filter out (ignore) those messages they are not interested in. One good point is that the network is simple to instal and there is no routing of traffic involved.

Now the bad news! The technology problems of using coaxial cable tapped in many points over long distances are NON-TRIVIAL. Work on micro-based implementations at the Rutherford lab. and at QMC has not been very encouraging - basically the network control is complex and requires a micro to handle it. Getting data into micros (and out) at 2-5 Mbaud is not cheap (or easy) and the software to handle it all is VERY complex. Reports are available from me about a low-cost ethernet we are evaluating at QMC. The XEROX ethernet has a raw data rate of 3Mbaud, as does the Rutherford ENET. The CNET at QMC is currently running at only 100 Kbaud. In other words, this is not an off-the-shelf network by any means and there is much work to be done to make it into one.

4. Ring Networks 8

Here all nodes in the network are connected in a ring. The connections, being point-to-point are easy and can use a wide variety of cable technology such as standard teletype twisted pairs, optical fibre, coax, etc. The ring network hardware at the Computer Laboratory in Cambridge is simple, replicable, easy to use and control. The data rates around the ring vary between 10 and 17 Mbaud. Rings are just as prone to failure as Ethernet coax is, however, but if that worries you, instal rudundant rings or many of them with gatweways in between. This is a promising candidate for off-the-shelf production after a little more work.

5. Other techniques

We discussed some other possibilities at the Workshop but I will leave Andy Hopper to describe those.

Conclusions

Local network technology for campus use in Universities has not got to the point yet where it cannot be called research. There is not a technology available which will meet the requirements outlined above **for a reasonable cost**. Work is proceeding at QMC, Cambridge, Edinburgh, Netcastle, Sussex(?), Warwick(?) on local network research and at QMC we are definately interested in using the fruits of our work on our own campus.

About this time next year there should be more encouraging news to report. Please would those who are developing networks of this nature make themselves known to me so that we can establish a forum for discussion of the issues as well as a descriptive workbook describing the nature of our interest.

A lot of the material for my part of the presentation came from Andy Hopper from Cambridge who has been very active in the field of rings and LSI design of universal switching modules. The friendly rivalry that exists between us on the question of technology and software has been a great source of stimulation to me and will, I hope, be of benefit to us all!

Thank you for making the workshop enjoyable,

Anthony West

A. Hopper
University of Cambridge
Computer Laboratory
Corn Exchange Street
Cambridge

LOCAL AREA COMPUTER NETWORKS - A BRIEF SURVEY

This paper gives an overview of the current state of local Area Computer Networks. A description of a number of local area computer network architectures is given, together with an outline of the Cambridge Computer Laboratory Ring System. Performance characteristics of some of the schemes are compared, and a brief indication of the design choices for a local network LSI chip is given.

LOCAL NETWORK ARCHITECTURES

Fig.1. shows the subdivision of local network technology. Recently the broadcast and ring schemes have been most popular. Such systems have been analysed in detail and their performance characteristics are similar, the delay being proportional to the number of nodes [HO78]. This is in contrast to the direct connection approach, or the crossbar switch, where the delay is constant, but the hardware rapidly becomes complicated as the number of nodes increases. There are a number of schemes which lie between these two extremes. The star network can offer low delays, but is dependent on a fast central switch, and is thus vulnerable to failure. It is also composed of dissimilar units and thus does not lend itself to implementation in LSI. The circuit switched approach utilises Nlog(N) interconnections, has worst case delay of log(N), but requires an external computer to set the switches to achieve the desired interconnection pattern [BE64,SU77]. Routing networks based on binary steering of packets between nodes have been proposed, but as yet none have been implemented in hardware.

Ring Systems

Rings allow a number of users to share a single communications link. As one of the most important parameters in network design is delay, ring systems are normally operated at low loads. This has the advantage that most of the irregularities inherent at high traffic levels do not occur.

Rings allow easy implementation of distributed switching and control functions and are thus well suited for use in distributed computer systems. There are also economic advantages to ring systems; they can be designed to be completely modular in nature and thus do not require a large initial investment.

The relatively poor reliability of a ring system poses a problem since the failure of any element will disable the entire network. This can be overcome by incorporating a parallel standby ring, and a number of such schemes have been proposed [ZA74]. Such

techniques are not suitable for rings with a small number of stations since the probability of failure of the reconfiguration units is greater than the probability of their improving network reliability. In any case ring failures will often be less serious than a breakdown in a centralised system in terms of fault location and repair times.

A ring network can be implemented in a number of ways. In the pre-allocated scheme a portion of the bandwidth is allocated to each user. This means quanta of service arrive at fixed times, but the system is inflexible to varying traffic loads. In the other ring systems the bandwidth available to each user changes with load, so that in a lightly loaded system one user can monopolise almost all the bandwidth. However, under heavy load conditions the bandwidth is shared equally and hogging by one source cannot occur. One way of achieving this is by using a permission token. The user can only transmit when in posession of the token, which is passed on downstream at the end of the packet. This means packets can be of variable length and line utilisation may be high. However, buffers must be provided, as well as hardware for detecting and restoring the token when it has been corrupted. The token can be physically implemented in a number of ways.

Another approach, called register insertion, is to place the packet in a shift register ready for transmission [W175]. The register is inserted into the data path at the appropriate time and the packet is shifted out and passes downstream. Such a system can be made to operate in either variable or fixed packet length modes. When used in the variable length mode the scheme for removing packets becomes complicated, especially under error conditions. In the fixed length mode the packet makes its way to the destination, and then back to the source where it is removed by taking the register out of the data path. This scheme has the property that some buffering of data is done in the network itself, however as the data path is frequently changed the probability of failure increases.

The Cambridge Computer Laboratory ring system is based on the empty slot principle. The empty slot system as originally proposed by Pierce suffers from hogging [P172]. This can be overcome if each packet makes a complete revolution of the ring and is cleared by the sender before being passed downstream. In addition, if the sender knows the number of packets in the ring, the packet can be cleared immediately when it returns. When the ring is first turned on, the slot structure must be created; in the Cambridge ring this task is performed by a monitor station, but it could be distributed among the other stations. The packet structure is shown in Fig.2. The start of packet (SOP) bit is followed by a bit which indicates whether the slot is full or empty. Next follows a bit which is used by the monitor station to delete packets which are circling indefinitely due to errors. Now follow four eight bit bytes used for destination and source addresses, and for data. There are two control bits at the end of the packet which are used for low level acknowledgements. These bits are set on the fly at the destination to indicate accepted, busy or unselected (rejected). If a packet returns with these bits unchanged it was not recognised at any destination. Each station posesses a select register which can be set to accept, or reject, all packets addressed to it, or to receive from one source only.

The ring is built using TTL logic and operates at 10 MHz, with a maximum distance of 200 meters between nodes. To improve reliability the logic for repeating the signal at each node is powered from the ring.

The Cambridge Ring is a fixed packet length system and thus an additional mechanism has to be provided to cater for variable length data. On the other hand data bytes can be transmitted

asynchronously, so that buffers may be dispensed with at the receiver as the reception of a message is suspendable at any point. Due to these design decisions, a direct memory access controller could be designed using redundant memory cycles and no storage. A disadvantage of most ring systems is that the delay through each node is at least one bit. In the Cambridge scheme this can be easily reduced, and thus a large number of nodes can be supported.

Broadcast Systems

Broadcast networks are characterised by the fact that a number of nodes may attempt to transmit at the same time, which results in all transmissions being corrupted. Such networks have been primarily developed to exploit the broadcast and multi-access capabilities of radio channels, although they can be implemented using a conventional cable system. The radio broadcast channel can be designed using satellite or ground radio. In the former case the round trip delay is approximately 0.27 seconds and in the latter the propagation delay is much smaller (microseconds). For a satellite channel, by the time a packet reaches the receiver, the transmission is past history and no control action can be taken. When the propagation delay between transmitter and receiver is small, the transmission can be aborted early on if this is desirable.

The broadcast network was first used in the Aloha system for connecting terminals to a central computer [AB70]. One channel was used for transmitting data to the computer, and another for transmitting acknowledgements back to the terminals. Each terminal transmits and stores one packet at a time and then initiates a time-out for the returning acknowledgement. The central computer can detect if a collision has taken place by using a suitable error check, in which case it does not send the acknowledgement. If no acknowledgement is received at the terminal before the time-out the packet is retransmitted. In order to avoid two terminals timing out and colliding repeatedly the time-outs for each are set to different values. An alternative scheme is for each terminal to choose the retransmission interval from a random distribution.

When the propagation delay between source and destination is small the Carrier Sense Multiple Access (Ethernet) system increases the theoretical channel capacity. Before transmission the channel is sensed, and if it is occupied the transmission is deferred until some time later. If the channel is sensed idle then the transmission proceeds, and is vulnerable to interference for a time equal to the propagation delay between the two most distant points in the system. If a collision takes place it is detected by both transmitters and the packets are retransmitted later. Once a transmission has been established it continues without interruption [ME76]. A number of protocols, some of which considerably increase throughput, have been proposed for the users action after sensing the channel [KL75]. However, like Aloha systems, the Ethernet system is characterised by the throughput tending to zero for large values of channel traffic.

Of all local network architectures the Ethernet approach has recently been most popular. This may not continue if the advantages of other network structures become better understood. Under low loads the Ethernet approach is advantageous as transmissions proceed with very little delay. However there are disadvantages to using Ethernet systems;

1. A portion of the bandwidth is wasted due to collisions.
2. As broadcast systems are unstable, additional hardware has to be provided to avoid the number of blocked terminals becoming arbitrarily large.

3. No low level acknowledgements are returned to the source.
4. As transmission speeds increase, a larger part of the packet will be lost before the transmission is aborted. The Ethernet approach is thus inappropriate at high transmission speeds.
5. A single break in the communication medium brings down the system due to unterminated signal reflections.
6. It is difficult to make the system synchronous.
7. Collision detection is difficult as the distant signal may be very weak.
8. Two way repeaters are required for stubs.
9. Packet buffers are required unless the interconnected machines are sufficiently complex.

Other Local Network Developments

There are a number of broadcast techniques which resolve the instability problem under overload conditions. These fall into two categories, dynamic control procedures and reservation schemes. Reservation schemes are suitable for systems in which each message is composed of several packets as the access request is made for complete messages. Dynamic control schemes require the user to take action to prevent channel saturation when the backlog of packets reaches a certain level.

Local network techniques have been used for linking very large machines. This generally entails the design of a sophisticated station which can communicate at data channel speeds. At the other extreme a number of microprocessor based systems are being built. These can be very cheap as in the QMC C-net [WE77], or suitable for applications where a large number of microprocessors are interconnected for control applications.

PERFORMANCE MEASURMENT

In this section the performance characteristics of various local network systems are compared. This is intended as a brief guide and a detailed analysis can be found in [HO78]. Fig.3. shows the delay for a number of ring systems against the source load. The delay is normalised in terms of the time for one bit to travel round an empty ring, and the source load in terms of ring transmission speed. The register insertion system is subdivided into systems with and without instant replacement of packets in the transmission shift register. It can be seen that the delay characteristics for all, except the pre-allocated, systems are similar; the slot scheme being slightly worse due to the presence of gap digits. At low loads the register insertion system is better because of the smaller non-alignment delays.

As the length of the ring increases the effect of the gap digits becomes smaller (for a given packet size). This is shown in Fig.4. where line utilisation is plotted against the number of stations. It is assumed each station contributes a fixed amount of electronic delay. As the system increases in size, the efficiency for the token and register insertion systems is unchanged, whereas for the slot system it tends to one. This is because for token and insertion the extra delay in the larger system is not utilised, whereas for slot this contributes to new slots and is in due course used for transmission. With the assumptions of the model the efficiency of the pre-allocated system tends to decrease with system size.

When comparing broadcast and ring systems it is found that at low loads the broadcast schemes have lower delays. This is because packets are not received back at the source and there are no non-alignment delays. However as the load increases clashes tend to occur more frequently and rings become better. For the broadcast schemes as the number of nodes increases, the maximum attainable throughput

decreases. In contrast ring systems either remain unchanged or improve.

A LOCAL NETWORK CHIP

Developments in silicon technology over the past fifteen years have mostly benefited CPU and memory hardware, while I/O and communications have lagged behind. The application areas of local networks are sufficiently broad to warrant the consideration of a station design in LSI.

If the hardware requirements of the different ring and broadcast systems are compared it is found that they are similar. This suggests that a single, general purpose chip design is possible, where the underlying hardware can be controlled in a number of ways to achieve the desired network structure. Such a chip is general purpose in nature, but is also very complex. At the other extreme a LSI implementation of simple ring scheme can be manufactured using cheap technology. The spectrum of chip complexity is shown in Fig.5., where the flexibility of the chip increases with the number of gates. Also shown, and intended as a rough guide, are the speeds and estimated costs of the different chip designs.

ACKNOWLEDGEMENT

The author would like to thank Professor D.J. Wheeler and Professor M.V. Wilkes for numerous discussions on the subject of local area computer networks.

References

AB70 Abramson, N. "Another alternative for computer communications", FJCC, 1970, pp. 695-702.

BE74 Benes V.E. "Optimal rearrangeable multistate connecting networks", BSTJ, Vol. 48, No.7, Jul. 1974, pp 1641-1656.

HO78 Hopper A. "Local area computer communication networks", Ph.D. Thesis, Computer Laboratory, University of Cambridge, (in preparation).

KL75 Kleinrock L., F.A.Tobagi, "Packet switching in radio channels", Parts 1,2, IEEE Trans on Comm., Vol. Com.23, No.12, Dec. 1975.

ME76 Metcalfe R.M., D.R.Boggs, "Ethernet: Distributed packet switching for local computer networks", Communications of the ACM, Vol.19, No.7, July 1976, pp.395-404.

PI72 Pierce J.R., "Network for block switching of data", BSTJ, Vol.51, No.6, July 1972, p.1133.

SU77 Sullivan H., T.R.Bashkow, "A large scale, homogeneous, fully distributed, parallel machine", Comp. Arch. News, Vol.5, No.7, March 1977, pp. 105-107.

WE77 West A.R., "A broadcast packet switched computer communications network", Design Progress Report, Queen Mary College, London.

WI75 Wilkes M.V., "Communication using a digital ring", Proc. of the Pacific Area Comp. Comm. Network System Symposium, Sendai, Japan, Aug. 1975, pp.47-56.

ZA74 Zafiropulo P. "Performance evaluation of reliability improvement techniques for single loop communications systems", IEEE Trans. on Comm., Vol. Com-22, No.6, June 1974, p.742.

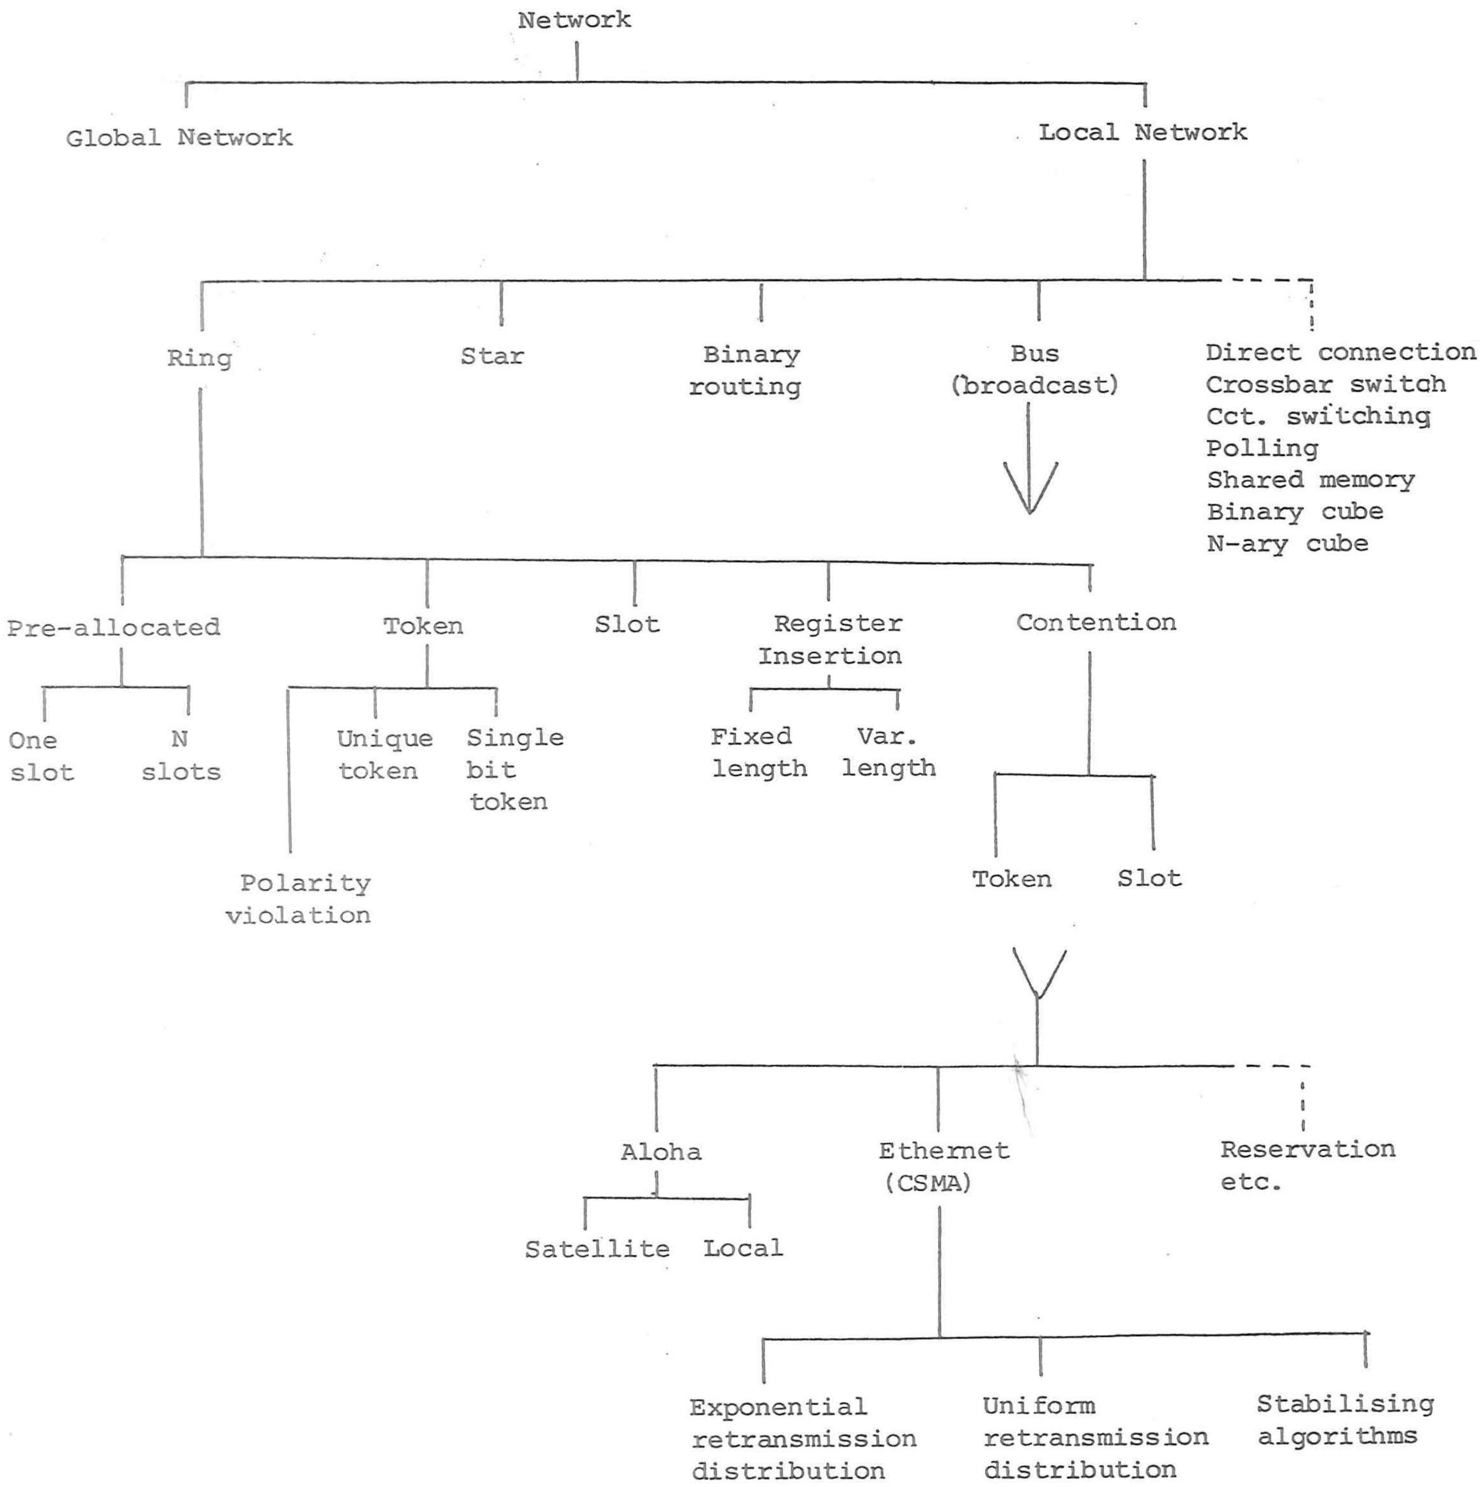

Fig. 1 Local Network Architectures

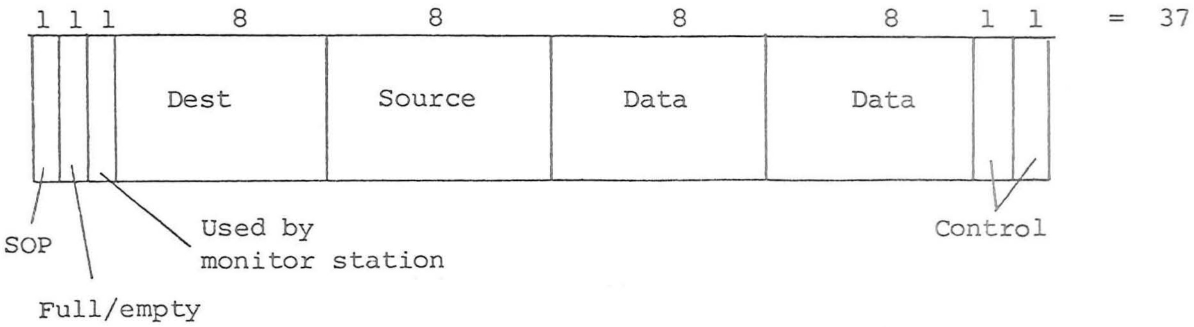

Fig. 2 Cambridge Ring Packet Format

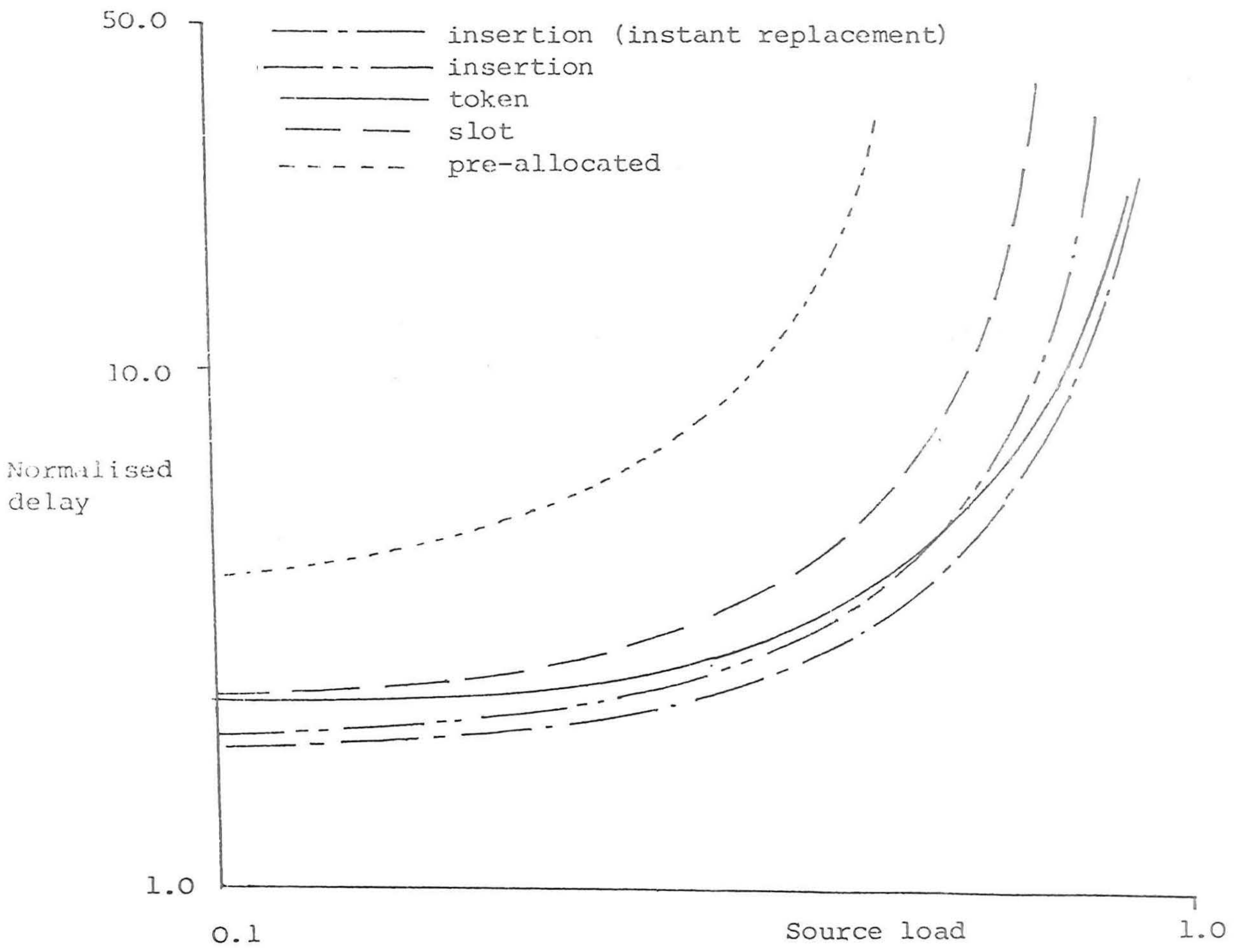

Fig. 3 Delays in Ring Networks

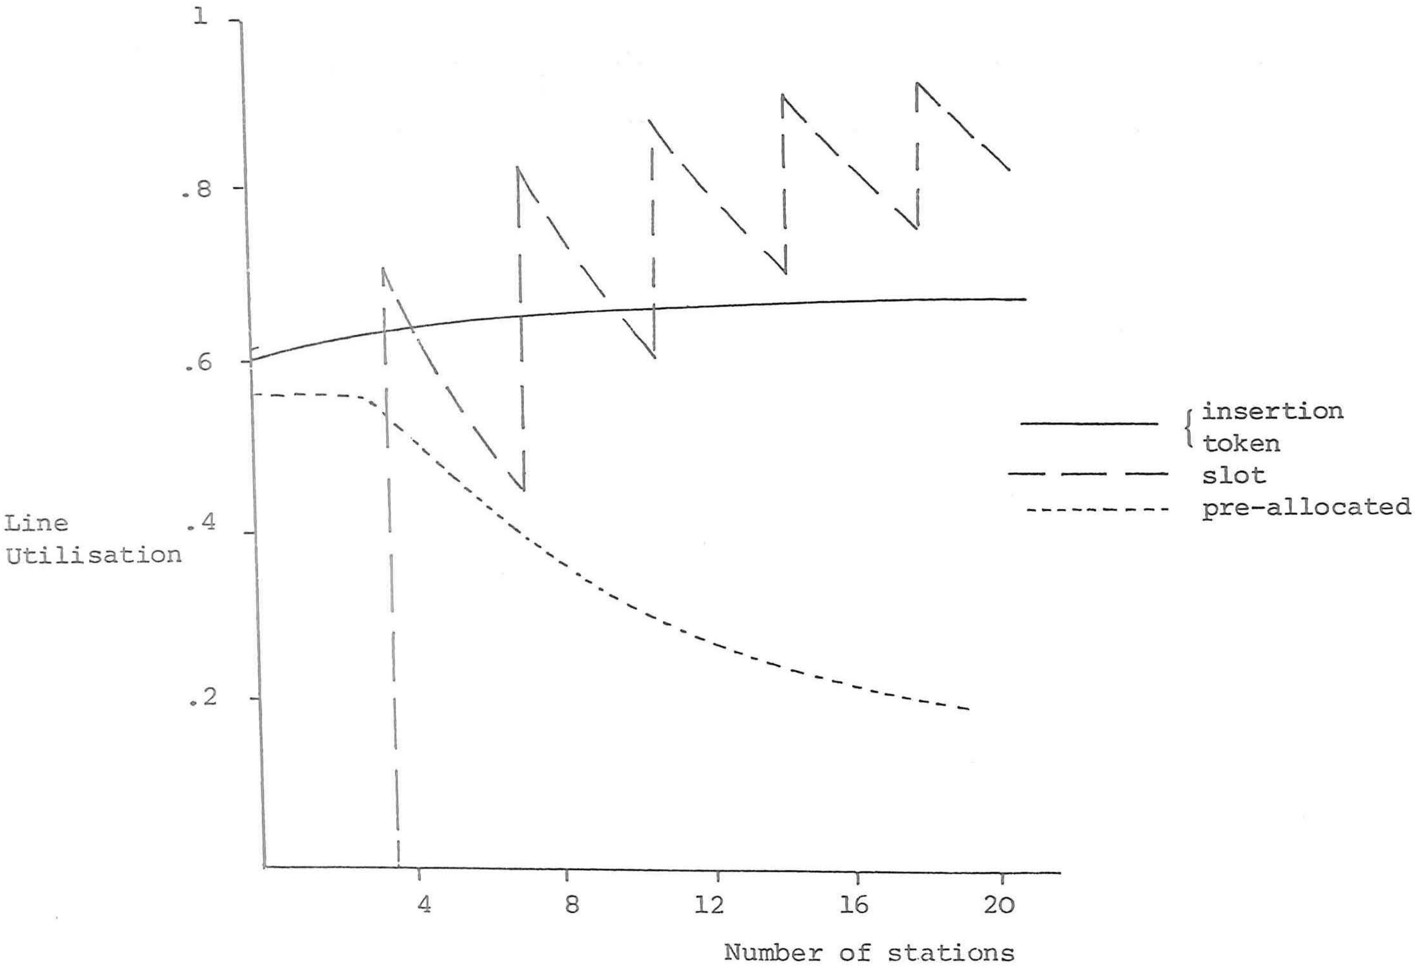

Fig. 4 Line Utilisation in Ring Networks

No. of Gates		<500	<700	>1500	Hybrid →
Networks implemented		-Binary routing -Simple slot scheme	-General purpose slot system	-Slot -Token -Ethernet -Aloha -Register insertion	
Logic family	ECL ←	ULA(CDI)	PMOS NMOS	SLF (I^2L) NMOS	
Speed		10 MHz	1 MHz	7 MHz	
Cost/chip		Cheap	Cheap	£50	
Investment		£5K	£15K	£50K	
Time scale		1 year	1-2 years	2-3 years	

Fig. 5 Chip Design Alternatives

GATEWAYS TO CAMPUS NETWORKS

A. J. Hinchley, C. J. Bennett

Dept. of Statistics & Computer Science

University College London

Our first consideration in building a campus gateway is in defining the functions that it ultimately performs for the user. Three major classes immediately come to mind.

i) The gateway may simple provide basic transport service for particular inter-process communications.

ii) The gateway may support terminal access into and out of the campus network system.

iii) The gateway may support some other particular service such as file or job transfer.

Of course the design and implementation of a gateway for any particular system may not result in all these services being supported, either implicitly or explicitly, and it can be seen from looking at the design of existing gateways (for example, those at the National Physical Laboratory and at University College London) that the facilities offered and the design followed can differ quite substantially. If we bear these general objectives in mind, however, a number of clear points emerge as design problems. It is these that will be the subject of this short talk.

General Design Issues

There are a number of general requirements that a good gateway design should meet. Firstly, the gateway should occupy a clearly defined position in the hierarchies of each of the networks it is connecting. Although the options of regarding the gateway as a packet switch, or as creating a new kind of network entity have been discussed, in practise we may assume that the gateway will appear as a host to

both networks. Certainly, for an X25 public network an X25 DTE connection is the only currently available way of attaching to the network.

This choice is not as clearcut as it might seem, as the gateway may be a separate machine or its function may be performed within an existing host system. Certainly if realised as a separate node processor we will want to apply to it all the same considerations we may have in mind for a network switching node - reliability, ability to control and reload in case of failure etc.

If it is implemented as a process within an existing host system these facilities may already exist in some form, although in many university environments these requirements will not be stringently met.

Secondly the gateway should be designed so that the network interconnection preserves local network autonomy. That is, existing network protocols and facilities should be available for local use in exactly the same way as they were before the network connection. Moreover, the gateway should act to protect each network against failures caused by actions in the other. It should not be possible, for instance, for a transnet file transfer to flood either the public net or the destination campus net with traffic to the point where one of these networks collapses, and it is the gateway's job to guard against this. Local network autonomy does not imply that the gateway must concentrate on its transnet role. There are facilities at the connection point that could be extremely useful to a connecting private network. We will elaborate on some of them, and possibly in the future we may find more attention paid to this area than so far.

Thirdly, it has to be decided whether the gateway is supporting end-to-end transnet protocols, or whether it is going to map one networks protocols into the next, thereby supporting a hop-by-hop communication. Hop-by-hop protocols will in general provide a poorer service than a good end-to-end service, as the protocol mappings are unlikely to be exact or complete. End-to-end service requires some fairly special conditions, however. There are only two situations in which it is reasonable to support it. The first occurs when the connected networks are either identical in structure or very close to it, in which case end-to-end protocols will emerge simply by extending the address space on each side to include the other. The second case occurs when the protocols concerned are network-independent. The file transfer protocol discussed yesterday is an example of these, and a great deal of research effort is going into this area. So far only a few complete protocols have emerged, and these have yet to meet with general acceptance though the file transfer protocol has made a promising start.

Finally, we have the very concrete design constraint in this case that all future campus network intercommunications will take place through X25 based regional or national networks. For the purposes of the rest of this discussion, we shall lump these nets into a single entity - the public network. This network will have a set of facilities and protocols which will be common for one side of all campus gateways connected to it.

SPECIFIC ISSUES

We will now consider some more specific issues in the light of the proceeding discussion.

Charging

The gateway will of course be responsible for all calls made outwards into the public network, and the bill will not distinguish source addresses within the private network unless mechanisms are specifically built in to the PTT charging such that sorting on the two available private digits is possible within the itemised billing information.

If the gateway is to attempt to maintain records of all calls itself it must be aware of the distance, packets passed, and any other information needed for charging. Ideally the public net might return the call cost at call clearing time, to the gateway. However as call charges will almost inevitably be calculated offline this is not feasible. The possibilities are that either a copy of the call record is in fact made available across the DTE interface at the end of the call or that a call charge code is returned in call accepted packets and the gateway then performs a packet count. The call charging will of course be reasonably simple for national calls. It is when international calls are made that charging complexities will creep in, but perhaps this is a problem that can be safely consigned to a low priority item on the designer's list.

If the gateway uses an existing switching node implementation then of course the charging information accumulated by the node software for PTT charging can be as well used for gateway calls. In fact it may be useful in any campus X25 network to implement the same call charging software as the PTT. The completed call records can then be channelled to a suitable database and although no-one would assume that the same order of billing would in fact be needed, nevertheless it provides a standard way of summing network statistics and performing more detailed charging analysis for detecting particular heavy users, or apportioning charges between several

participating authorities within the campus network.

Access Control

How can we provide control on outgoing calls? Do we block on certain destination DTE addresses (international for example), or can we block on source local net address. Can we obtain in any way the originating subscriber ID and make the check on that? Of course if the gateway is nearer to a host service, we may be able to use the checking facilities of the host:- login etc to perform some control. However performing such control on terminal connections for instance immediately separates out this particular service as having certain procedures which may not otherwise be appropriate, for instance, in the context of a gateway supporting end-to-end protocols and providing only a basic transport service.

Technical issues

There are a large number of technical issues connected with gateways. We shall only examine two of these - the provision of addressing, and the support of specific services - in this talk. We have already mentioned flow control as a problem when discussing local network autonomy; other problem areas include routing, which will occur when a networking is connected to the public net at more than one point; fragmentation and reassembly of packets, which occurs when the two networks have different packet sizes; and error detection and reporting. We can examine addressing and service support in the light of a very specific design constraint in the case of British University campus connections, which fall into two groups:

-campus X25 public X25

-campus proprietary protocols-public X25

When we examine the two groups in more detail we discover that the difference in approach is sufficient that we may discuss the two groups completely separately.

The X25 Local net

We may assume that in a private network, most if not all of the conventions adopted in attaching to private nets will be retained. After all the attraction of having a campus X25 network will be to retain the same standard software modules possibly already developed to support public net access. We may expect standard level 3 and level 2 access, PAD or virtual terminal access, together with further higher-level protocol standards that may be developed.

This enables us then to propose an X25 DTE/DTE gateway able to concatenate virtual calls and carry any higher level protocols transparently. The gateway will in fact be very much like an X25 switching node, but providing DTE access rather than DCE. In fact the network unit node switch specification rather implies these sorts of functions.

Addressing

We would like to have an addressing scheme that was good for all subscribers reachable from the X25 public network(s) irrespective of whether subscribers are directly attached to the X25 network, or attached to some connecting private or campus network. Unfortunately no recommendations exist to make this possible. Datapac in Canada have allowed two final digits on each subscriber address which can obviously address up to 100 entities in a private net. This will be satisfactory for many situations. It does not of course allow much scope for any hierarchical addressing within the private net of nets.

An alternative is to use the data field for some specific sub-addressing. The drawback here is that it is very difficult to standardise and of course one is requiring the source DTE or PAD to encode a private address in a very net-specific way.

We can of course map the external address into a part of a much larger internal address set. This will be probably sufficient for a rather definable subset of internal subscribers who wish to get

internet access or receive internet incoming calls. It does however cause extra disturbance where the internal set is changing forcing the mapping into a non-linear untidy grouping.

Call handling

The nature of the virtual call on either side is clearly conforming to an X25 DTE specification and the requirement is to provide effectively a concatenated call facility.

-we must provide call tables on a per connection basis

-we must decide whether to cascade the call request, or to accept the call at the gateway before trying the private net destination

-we must decide what features have end-to-end significance rather than local significance.

Again all the same things as a network switch providing concentrated call service.

The non-X25 local net

Here, existing gateways between networks provide an example of the problems and solutions. Because protocols on either side of the gateway will be different we must provide a level of protocol mapping at the gateway. As soon as we map protocols we may need to be selective about what services can be supported. Only a perfect mapping of the lowest level of service would for instance allow us to support all higher level services (always assuming that higher level services have been carefully structured on top of the low-level service). The higher level protocols may not of course even be the same, within the private net as opposed to in the public network. Here we must choose from a number of alternatives:

-implement additional high-level protocols in the local net such that when going internet, the compatible alternate HLPs are used instead of the regular campus net facilities.

-attempt to map each higher-level protocol separately at the gateway.

-designate a host in the campus net which can provide internet services.

We may conclude that if at all possible, an automatic mapping of addresses is preferable, and the EPSS process labels is an obvious current way in which this problem can be solved. The need to standardise on the two X25 DTE digits or some other acceptable field description is underlined then in this type of gateway as well if we wish to achieve some level of internet connection which is analogous in its simpicity of operation to making a local campus connection.

Call-level mapping

The local network may not have the concept of a virtual call. It is however implicit in mapping into an X25 public net, that the appearance of a virtual call must be constructed in the campus net, if an end-to-end virtual call service were to be offered. The alternative is of course to make use of an existing host connected to both networks, and then use the logging-in procedures to support incoming terminal traffic, and provide a named service which is in fact an outgoing terminal connection to the other network. The Electric terminal support system at Rutherford is an obvious indigenous example where this is achieved. In the US Telenet users may log in to Tenex hosts from the Telenet network, and use a TELNET service to make an outgoing ARPANET connection.

We may conclude that a direct call-level mapping is only possible if direct address mapping can be supported, and also where end-to-end higher-level protocols are compatible across the networks.

Terminal-level mapping

We have already quoted the UCL SWITCH gateway as an example of this mapping level. We may use the experience gained here to say something about the problems. If we used a parameterised

terminal protocol such as EPSS VPT or the PAD, we may find it very difficult to try and map all parameters across the gateway. For the UCL EPSS/ARPANET gateway for instance we support a selected set of terminals from EPSS and map into the ARPANET virtual terminal. In the reverse direction, it is simpler as there can be a fixed mapping into the EPSS VPT. Problems arise in the following areas:-

-the handling of control characters

-different conventions in message as opposed to character

-reconciling local and remote echoing features

-supporting sophisticated terminals across the gateway.

It can be seen that providing a general service is extremely difficult and a decision must be taken on very local conditions on whether the resulting service is satisfactory for the target users.

File transfer mapping

Here we may attempt direct mapping of the protocol at the gateway or else use some staged system to operate more as a batch-type service. Here the service system might in one stage get the file from the remote network destination, and in a second stage deliver it to the user:- different file transfer protocol arrangements pertaining in each transfer of course.

A direct mapping is difficult given the parameterised nature of FTP set-up exchanges. The result is inevitably a rather cut-down subset of what a user might rightly expect to have available as a service. The planned UCL ARPA/EPSS file transfer will come into this category, and as a longer-term plan it is infinitely preferable to explore a true end-to-end service, or to examine the benefits of a staged service.

To summarise this area a mapping at the lowest level is

preferable if at all possible. Mapping at the higher level is rarely totally satisfactory, and must be done for each level of protocol. The use of a service host is probably a better solution for that category of situations.

OVERALL SUMMARY

We have tried to present the possible solutions in very distinct categories. Inevitably in practice the distinctions will blur. However a distinct polarization is very clear. Where a call-level mapping can be achieved the gateway will look very much like a switching node in its construction, maintenance- and operation. Where this is not possible the gateway is most likely to be a function on a service host system.

Short list of useful references on gateways

non-X25 related

Interconnection of packet switching networks:theory and practice. M.Gien,J.Laws,R.Scantlebury EUROCOMP75

Problems of connecting networks with a gateway computer.
P.L.Higginson,A.J.Hinchley EUROCOMP75

Interconnection of computer networks C.Sunshine Computer Networks voll. no.1

Initial experiences with an EPSS service.
P.L.Higginson, Z.Fisher EUROCOMP78 (to be published in May

RIG, Rochester's Intelligent gateway: system overview.
J.Ball,J.Feldman,J.Low,R.Rashid,P.Rovner.
IEEE Trans. Soft. Engineering Dec. 76

X2 related

Interconnection of networks. L.Pouzin
Infotech report on future networks (to be published July)

An example of connecting aX25 based computer network to a datagram based computer network. F.Vogt EUROCOMP78

International Interconnection of packet switched networks
L.Roberts p239 ICCC76

The extension of network services across many networks.
A.J.Hinchley,P.Kirstein EUROCOMP78

Interconnection of international X25 networks.
G.Grossman,A.J.Hinchley (to be submitted for USA/Japan conference September 1978)

GATEWAY FUNCTIONS

- BASIC TRANSPORT SERVICE
- TERMINAL SUPPORT
- FILE AND JOB TRANSFER SUPPORT

GENERAL DESIGN ISSUES

- LEVEL OF CONNECTION
- LOCAL NETWORK AUTONOMY
- END-TO-END VS HOP-BY-HOP
- X25 PUBLIC NET

SPECIFIC ISSUES

- CHARGING
- ACCESS CONTROL
- X25 LOCAL NET VS OTHER
  - . END - TO - END ADDRESSING
  - . SERVICE SUPPORT
  - . ERROR HANDLING

X25 LOCAL NET

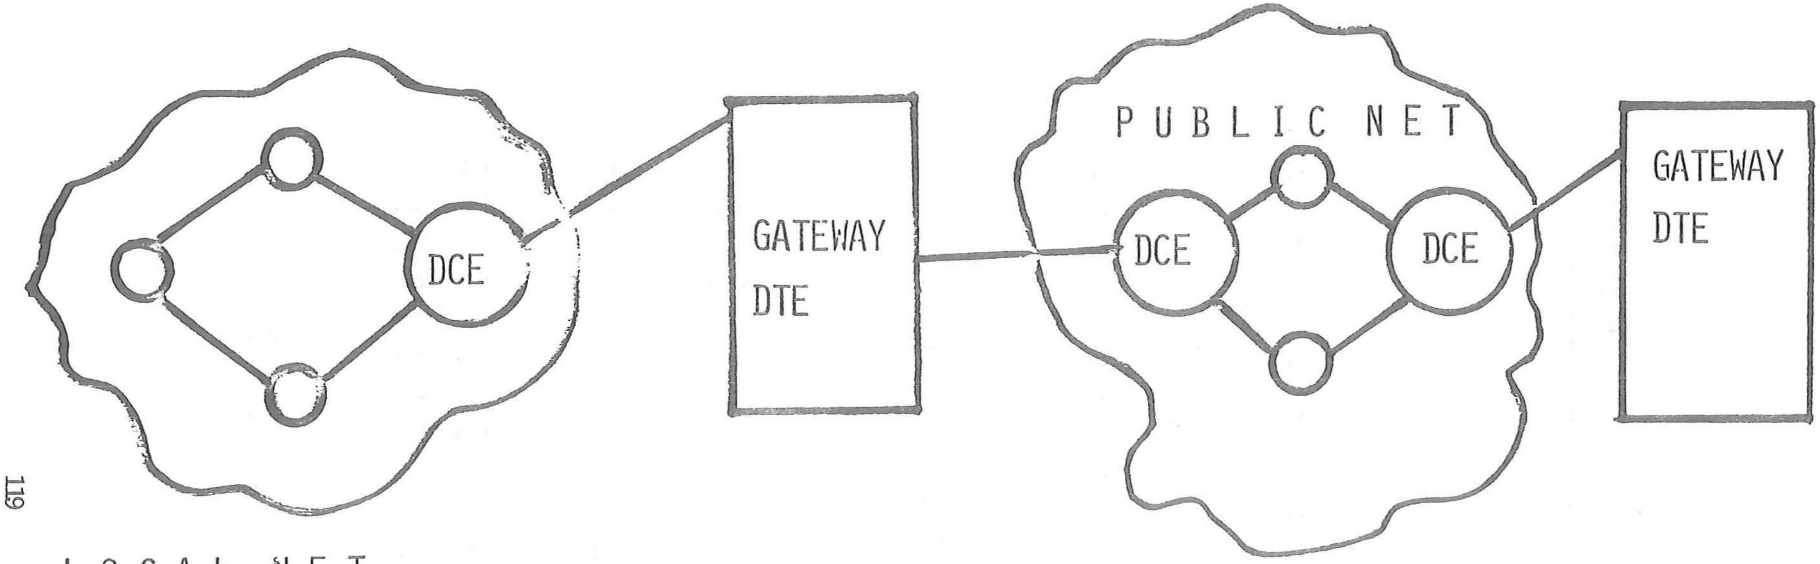

LOCAL NET

X25 LOCAL NET

- ADDRESSING

  . DIGITS IN SUBSCRIBER ADDRESS

  . DATA FIELD

  . ADDRESS MAPPING

- CALL HANDLING

  . CONCATENATION ON PER CONNECTION BASIS

  . CASCADE OR STEP-BY-STEP

  . SEPARATE END-TO-END FROM LOCAL FEATURES

= ACCESS CONTROL

NON-X25 LOCAL NET

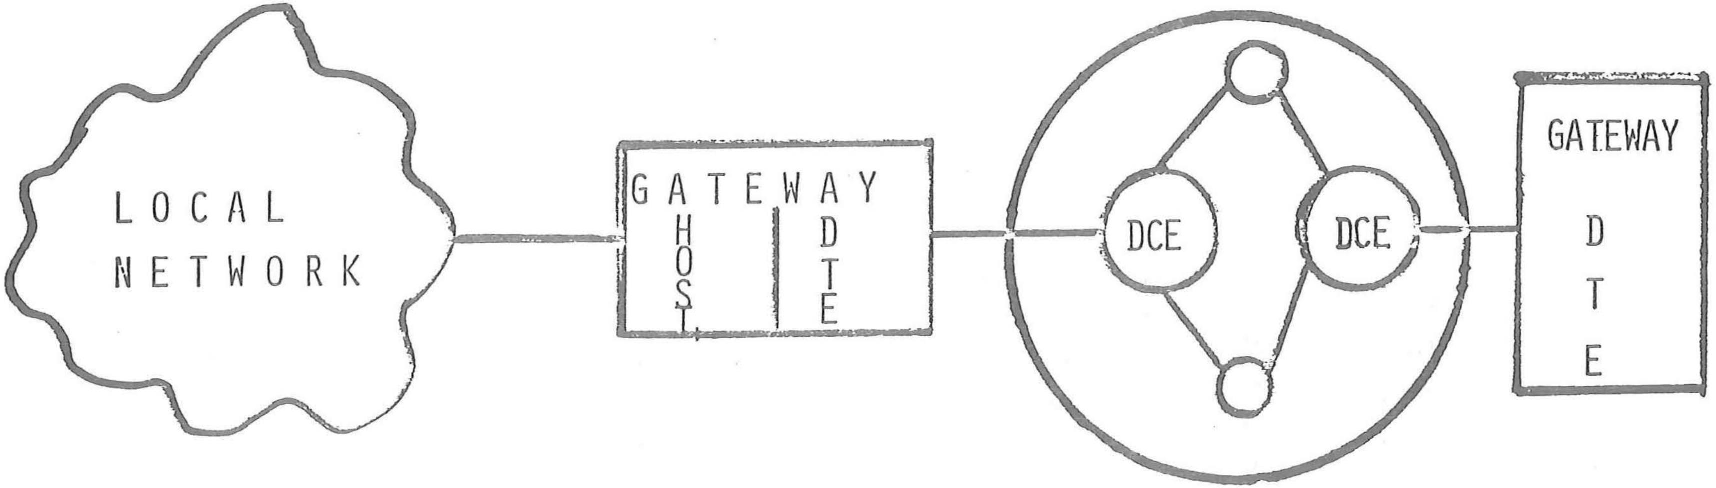

PUBLIC NET

NON-X25 LOCAL NET

- PROTOCOL MAPPING
  - ALTERNATE PROTOCOLS
  - MAPPINGS AT GATEWAY
  - INTERNET SERVICE SITE
- ADDRESSING
  - MAPPING
  - TERMINAL ACCESS
- CALL HANDLING
  - VIRTUAL CALL

- TERMINAL MAPPING
  - CONTROL CHARACTERS
  - MESSAGE VS CHARACTER PROTOCOLS
  - LOCAL VS REMOTE ECHO
  - SOPHISTICATED TERMINALS
- FILE TRANSFER
  - DIRECT MAPPING
  - STAGED SERVICE

CONCLUSIONS AND FUTURE TASKS

Chairman: R. Rosner

Speaker: M.B. Williams

Purpose of Networkshops

The main purpose of holding networkshops has changed from that originally envisaged. After the first workshop it was intended that various individuals should act as rapporteurs to collect information about particular subjects, form "task forces" where necessary and report back on progress at the second workshop. Several of the rapporteurs identified have been involved in national and international study groups and the original concept has not been followed through. The progress of such study groups cannot be controlled by the University community but such university and research council involvement is of high priority.

The function of these Networkshops has therefore evolved to a concern with:

a) Information - enabling the participants to find out what was happening nationally and internationally from some of the experts in different fields,

b) Progress - monitoring the progress of university plans and national and international standards,

c) Needs - identifying the needs of the University & Research Councils community.

To ensure that the information disseminated at Networkshop 2 was not lost a rapidly produced comprehensive report of the proceedings was essential. It was agreed that a further Networkshop was required and this would be held at Cambridge on 26-27 September. An organising Committee was formed, comprising Peter Linington and others from Cambridge, the Network Unit, and John Rice from Liverpool. By the third Networkshop it is believed that significant progress will have been made in the formulation of plans by the Computer Board, the Science Research Council and the Post Office.

A number of Specific issues arose at the Workshop and detailed discussion followed of the topics to be covered at Networkshop 3.

Specific Issues

A. Method of Access to X25 Nets.

The method of access to X25 was effectively going to be decided/influenced by four bodies:

1. Study Group 4A
2. National University & Research Councils Centres
3. P.O. Provision of PSE
4. Supply of Local Area (Campus) PSE's and discussions with prospective manufacturers.

A meeting was to take place in the next few weeks between the Network Unit and the National Centres to discuss X25 access. The P.O. should have a significant number of statements to make regarding documentation, tariff structures and facilities required for Campus and regional switches by the next workshop in September. The Network Unit will be approaching the P.O. in the very near future to get some clarification of P.O. supplied Campus switches with regard to pricing, facilities and method of access.

The main points regarding X25 were that the Universities should be ready to convert to the P.O. X25 net by 1.1.1980 and that the interfacing of machines to X25 should only be undertaken once per particular model. As the Computer Board has written a letter to the manufacturers stating that they would only purchase machines that could connect to X25, a quick resolution of method of access was required to enable centres to talk sensibly to them about this requirement.

The next workshop will review this problem.

Roland Rosner is encouraging the establishment of groups of users with the same machine range in an attempt to present the duplication of work and the Network Unit undertook to prepare a review of the machine ranges for which progress to X25 interfaces is known to be currently taking place or is planned.

B. High Level Protocols

The chances of any agreements appearing from national or international standards bodies within the timescales to which the Universities were working seemed remote. It was felt that the Universities would probably have to take some unilateral action over the adoption of standards for high level protocols.

(i) FTP The creation of an FTP user group was suggested, but no chairman was proposed. The authors of the FTP document have thought of forming such a group and any prospective implementors should fill in the questionnaire provided with the document. The FTP standard was felt now to have been developed to a stage where it could be adopted by the university and research councils community.

(ii) Transport Stations
There seemed little likelihood of any standard emerging quickly. Mervyn Williams had approached the Department of Industry to see if this area could be identified as a project for national funding and was hopeful that something might come of this. General concern was expressed over the time-scales of such a project.

(iii) VTP The need for co-ordinated action was expressed.

(iv) JTP No standards body appeared to be working on JTP at present and it was felt that the community represented at the workshop would need to come to a private agreement over a standard to adopt.

C. P.O. Monopoloy

The P.O. had written to the Network Unit over the position of Regional Networks and the Network Unit will be replying shortly after obtaining approval from CB and SRC. A lot of support had been shown for P.O. operated P.S.E., but the crucial factor was price.

D. Local Area/Campus Networks

Considerable interest had been shown in Campus Networks. It was felt that an information exchange was required and a mechanism set up to keep people informed. A day long seminar would be arranged by Tony West and the Network Unit to start this off. The problem with Campus Nets was that they fell on the boundary between Research & Service and it was difficult to obtain funding for development projects. The Network Unit thought that there was a mechanism to obtain money from SRC for such purposes however.

A number of Universities have the same type of mainframes and some savings could be made in some form of common development. 1900 mainframes were mentioned particularly.

E. Hardware

The hardware field was developing rapidly and it was agreed that a further session at a future Networkshop would be needed.

Good information on what was available and a pooling of experience is required.

A new serial interface specification (#449), effectively a composite of V24 and V35, has recently been produced for forwarding to CCITT. Edinburgh has copies and these may be obtained from Ray Chisholm.

F. International Networks

It was felt that we needed to be better informed of the services to be offered by EURONET, EIN etc. Peter Higginson had copies of a new EURONET glossy which he would pass to R. Rosner for distribution.

Conclusion

At least 6 specialist areas for consideration at Networkshop 3 have been identified.